A Preconditioned Iterative Interior Point Approach to the Conic Bundle Subproblem

Christoph Helmberg*

October 21, 2021

Abstract

The conic bundle implementation of the spectral bundle method for large scale semidefinite programming solves in each iteration a semidefinite quadratic subproblem by an interior point approach. For larger cutting model sizes the limiting operation is collecting and factorizing a Schur complement of the primal-dual KKT system. We explore possibilities to improve on this by an iterative approach that exploits structural low rank properties. Two preconditioning approaches are proposed and analyzed. Both might be of interest for rank structured positive definite systems in general. The first employs projections onto random subspaces, the second projects onto a subspace that is chosen deterministically based on structural interior point properties. For both approaches theoretic bounds are derived for the associated condition number. In the instances tested the deterministic preconditioner provides surprisingly efficient control on the actual condition number. The results suggest that for large scale instances the iterative solver is usually the better choice if precision requirements are moderate or if the size of the Schur complemented system clearly exceeds the active dimension within the subspace giving rise to the cutting model of the bundle method.

Keywords: low rank preconditioner, quadratic semidefinite programming, nonsmooth optimization, interior point method
MSC 2020: 90C22, 65F08; 90C06, 90C25, 90C20, 65K05

1 Introduction

In semidefinite programming the ever increasing number of applications [3, 32, 6, 20] leads to a corresponding increase in demand for reliable and efficient solvers. In general, interior point methods are the method of choice. Yet, if the order of the matrices gets large and the affine matrix functions involved do not allow to use decomposition or factorization approaches such as proposed in [24, 4, 7], general interior point methods are no longer applicable. The limiting factors are typically memory requirements and computation times connected with forming and factorizing a Schur complemented system matrix of the interior point KKT system. In this case, the spectral bundle method of [19] offers a viable alternative.

The spectral bundle method reformulates the semidefiniteness condition via a penalty term on the extremal eigenvalues of a corresponding affine matrix function and assumes these eigenvalues to be efficiently computable by iterative methods. In each step it selects a subspace close to the current active eigenspace. Then the next candidate point is determined as the proximal point with respect to the extremal eigenvalues of the affine matrix function projected onto this subspace. The proximal point is the optimal solution to a quadratic semidefinite subproblem whose matrix variable is of the order of the dimension of the approximating subspace. If the subspace is kept small, this allows to find approximately optimal solutions in reasonable time. In order to reach solutions of higher precision it seems unavoidable to go beyond the full active eigenspace

*Department of Mathematics, Chemnitz University of Technology, D-09107 Chemnitz, Germany, helmberg@mathematik.tu-chemnitz.de, www.mathematik.tu-chemnitz.de/∼helmberg

[17, 9]. In the current implementation within the callable library ConicBundle [15] the quadratic subproblem is solved by an interior point approach. Again for each of its KKT systems the limiting work consists in collecting and factorizing a Schur complement matrix whose order is typically the square of the dimension of the active eigenspace. The main question addressed here is whether it is possible to stretch these limits by developing a suitably preconditioned iterative solver that allows to circumvent the collection and factorization of this Schur complement. The focus is thus not on the spectral bundle method itself but on solving KKT systems of related quadratic semidefinite programs by iterative methods. Even though the methodology will be developed and discussed for low rank properties that arise in the ConicBundle setting, some of the considerations and ideas should be transferable to general conic quadratic optimization problems whose quadratic term consists of a positive diagonal plus a low rank Gram matrix or maybe even to general positive definite systems of this form.

Here is an outline of the paper and its main contributions. Section 2 provides the necessary background on the bundle philosophy underlying ConicBundle and derives the KKT system of the bundle subproblem. The core of the work is presented in Section 3 on low rank preconditioning for a Gram-matrix plus positive diagonal. For slightly greater generality, denote the cone of positive (semi)definite matrices of order m by $\mathbb{S}^m_{++}$ ($\mathbb{S}^m_+$) and let the system matrix be given in the form

$$H = D + VV^\top \qquad \text{with } D \in \mathbb{S}^m_{++}, V \in \mathbb{R}^{m\times n},$$

where it is tacitly assumed that D^{-1} times vector and V times vector are efficiently computable. Typically $n \leq m$ but whenever n is sizable one would like to approximate V by a matrix $\hat{V} \in \mathbb{R}^{m\times k}$ with significantly smaller $k < n$ to obtain a preconditioner $\hat{H} = D + \hat{V}\hat{V}^\top$ with low rank inverse $\hat{H}^{-1} = D^{-1} - D^{-1}\hat{V}(I_k + \hat{V}^\top D^{-1}\hat{V})^{-1}\hat{V}^\top D^{-1}$. Comparing this to the inverse of H, the goal is to capture the large eigenvalues of $V^\top D^{-1}V$, more precisely the directions belonging to large singular values of $D^{-\frac{1}{2}}V$. By the singular value decomposition (SVD) this can be achieved by the projection onto a subspace, say $D^{-\frac{1}{2}}VP$ for a suitably chosen orthogonal $P \in \mathbb{R}^{n\times k}$. Because the full SVD is computationally too expensive, two other approaches will be developed and analyzed here. In the first, Section 3.1, the orthogonal P is generated by a Gaussian matrix $\Omega \in \mathbb{R}^{n\times k}$. In the second, Section 3.2, some knowledge about the interior point method leading to V will be exploited in order to form P deterministically.

The projection onto a random subspace may be motivated geometrically by interpreting the Gram matrix $VV^\top$ as the inner products of the row vectors of V. The result of Johnson-Lindenstrauss, cf. [1, 8], allows to approximate this with low distortion by a projection onto a low dimensional subspace. In matrix approximations this idea seems to have first appeared in [27]. In connection with preconditioning a recent probabilistic approach is described in [21] in the context of controlling the error of a LU preconditioner. [13] gives an excellent introduction to probabilistic algorithms for constructing approximate matrix decompositions and provides useful bounds. Building directly on their techniques we provide deterministic and probabilistic bounds on the condition number of the random subspace preconditioned system in theorems 6 and 7. In comparison to the moment analysis of the Ritz values of the preconditioned matrix presented in Theorem 4, the bounds seem to fall below expectation and are maybe still improvable. Random projections do not require any problem specific structural insights, but it remains open how to choose the subspace dimension in order to obtain an efficient preconditioner.

In contrast, identifying the correct subspace seems to work well for the deterministic preconditioning routine. It exploits structural properties of the KKT system's origin in interior point methods. Within interior point methods iterative approaches have been investigated in quite a number of works, in conjunction with semidefinite optimization see *e.g.* [31, 25]. These methods were mostly designed for exploiting sparsity rather than low rank structure. During the last months of this work an approach closely related to ours appeared in [12]. It significantly extends ideas of [33] for a deterministic preconditioning variant. It assumes the rank of the optimal solution to be known in advance and provides a detailed analysis for this case. Their ideas and arguments heavily influenced the condition number analysis of our approach presented in theorems 2 and 9. In contrast to [12], our algorithmic approach does not require any a priori knowledge on the rank

of the optimal solution. Rather, Theorem 9 and Lemma 12 motivate an estimate on the singular value induced by certain directions associated with active interior point variables, that seems to offer a good indicator for the relevance of the corresponding subspace.

In Section 4 the performance of the preconditioning approaches is illustrated relative to the direct solver on sequences of KKT systems that arise in solving three large scale instances within ConicBundle. The deterministic approach turns out to be surprisingly effective in identifying a suitable subspace. It provides good control on the condition number and reduces the number of matrix vector multiplications significantly. The selected instances are also intended to demonstrate the differences in the potential of the methods depending on the problem characteristics. Roughly, the direct solver is mainly attractive if the model is tiny, if significant parts of the Schur complement can be precomputed for all KKT systems of the same subproblem or if precision requirements get exceedingly high with the entire bundle model being strongly active. In general, however, the iterative approach with deteriministic preconditioner can be expected to lead to significant savings in computation time in large scale applications. In Section 5 the paper ends with some concluding remarks.

Notation. For matrices or vectors $A, B \in \mathbb{R}^{m\times n}$ the (trace) inner product is denoted by $\langle A, B\rangle = \operatorname{tr} B^\top A = \sum_{ij} A_{ij}B_{ij}$. $A \circ B = (A_{ij}B_{ij})$ denotes the elementwise or Hadamard product. Consider symmetric matrices $A, B \in \mathbb{S}^n$ of order n. For representing these as vectors, the operator $\operatorname{svec} A = (A_{11}, \sqrt{2}A_{21}, \ldots, \sqrt{2}A_{n1}, A_{22}, \sqrt{2}A_{32}, \ldots, A_{nn})^\top$ stacks the columns of the lower triangle with offdiagonal elements multiplied by $\sqrt{2}$ so that $\langle A, B\rangle = \operatorname{svec}(A)^\top \operatorname{svec}(B)$. For matrices $F, G \in \mathbb{R}^{k\times n}$ the symmetric Kronecker product $\otimes_s$ is defined by $(F\otimes_s G)\operatorname{svec}(A) = \frac{1}{2}\operatorname{svec}(FAG^\top + GAF^\top)$. The Loewner partial order $A \succeq B$ ($A \succ B$) refers to $A - B \in \mathbb{S}^n_+$ ($A - B \in \mathbb{S}^n_{++}$) being positive semidefinite (positive definite). The eigenvalues of A are denoted by $\lambda_{\max}(A) = \lambda_1(A) \geq \cdots \geq \lambda_n(A) = \lambda_{\min}(A)$. The norm $\|\cdot\|$ refers to the Euclidean norm for vectors and to the spectral norm for matrices. I_n (I) denotes the identity matrix of order n (or of appropriate size), the canonical unit vectors e_i refer to the i-th column of I. Unless stated explicitly otherwise, $\mathbb{1}$ denotes the vector of all ones of appropriate size. $\mathbb{E}$ refers to the expected value of a random variable, Var to its variance and $\mathcal{N}(\mu, \sigma^2)$ to the normal or Gaussian distribution with mean μ and standard deviation σ.

2 The KKT system of the ConicBundle Subproblem

The general setting of bundle methods deals with minimizing a (mostly closed) convex function $f\colon \mathbb{R}^m \to \overline{\mathbb{R}} := \mathbb{R}\cup\{\infty\}$ over a closed convex ground set $C \subseteq \operatorname{dom} f$ of simple structure like $\mathbb{R}^m$, a box or a polyhedron,

$$\text{minimize } f(y) \quad \text{subject to } \ y \in C.$$

Typically, f is given by a first order oracle, i. e., a routine that returns for a given $\bar{y} \in C$ the function value $f(\bar{y})$ and an arbitrary subgradient $g \in \partial f(\bar{y})$ from the subdifferential of f in $\bar{y}$. Value $f(\bar{y})$ and subgradient g give rise to a supporting hyperplane to the epigraph of f in $\bar{y}$. The algorithm collects these affine minorants in the *bundle* to form a cutting model of f. It will be convenient to arrange the value at zero and the gradient in a pair $\omega = (\gamma = f(\bar{y}) - \langle g, \bar{y}\rangle, g)$ and to denote, for $y \in \mathbb{R}^m$, the minorant's value in y by $\omega(y) := \gamma + \langle g, y\rangle$.

Let $\mathcal{W}_f = \{\omega = (\gamma, g) \in \mathbb{R}^{1+m}\colon \gamma + \langle g, y\rangle \leq f(y), y \in \mathbb{R}^m\}$ denote the set of all affine minorants of f. For closed f we have $f(y) = \sup_{\omega\in\mathcal{W}_f} \omega(y)$. Any compact subset $W \subseteq \mathcal{W}_f$ gives rise to a minorizing cutting model of f,

$$W(y) := \max_{\omega\in W} \omega(y) \leq f(y), \quad y \in \mathbb{R}^m.$$

At the beginning of iteration $k = 0, 1, \ldots$ the bundle method's state is described by a current *stability center* $\hat{y}_k \in \mathbb{R}^m$, a compact cutting model $W_k \subseteq \mathcal{W}_f$, and a proximity term, here the square of a norm $\|\cdot\|^2_{H_k} := \langle \cdot, H_k \cdot\rangle$ with positive definite H_k. The method determines the next candidate $y_{k+1} \in \mathbb{R}^m$ as minimizer of the *augmented model* or *bundle subproblem*

$$y_{k+1} = \operatorname*{argmin}_{y\in C} W_k(y) + \tfrac{1}{2}\|y - \hat{y}_k\|^2_{H_k}. \tag{1}$$

Solving this bundle subproblem may be viewed as determining a saddle point $(y_{k+1}, \bar{\omega}_{k+1} = (\bar{\gamma}_{k+1}, \bar{g}_{k+1})) \in C \times \operatorname{conv} W_k$, which exists for any closed convex C by [29], Theorems 37.3 and 37.6, due to the strong convexity in y and the compactness of W_k,

$$\begin{aligned} \bar{\omega}_{k+1}(y_{k+1}) + \tfrac{1}{2}\|y_{k+1} - \hat{y}_k\|^2_{H_k} &= \inf_{y \in C} \sup_{\omega=(\gamma,g)\in W_k} \gamma + \langle g, y\rangle + \tfrac{1}{2}\|y - \hat{y}_k\|^2_{H_k} \\ &= \sup_{\omega \in \operatorname{conv} W_k} \inf_{y \in C} \omega(y) + \tfrac{1}{2}\|y - \hat{y}_k\|^2_{H_k}. \end{aligned}$$

Strong convexity in y ensures uniqueness of y_{k+1}. First order optimality with respect to y implies

$$\bar{g}_{k+1} + H_k(y_{k+1} - \hat{y}_k) \in -N_C(y_{k+1}),$$

where $N_C(y)$ denotes the normal cone to C at $y \in C$. In the unconstrained case of $C = \mathbb{R}^m$ the *aggregate* $\bar{\omega}_{k+1}$ is also unique. Whether unique or not, the *aggregate* will refer to the solution $\bar{\omega}_{k+1} \in \operatorname{conv} W_k$ produced by the algorithmic approach for solving (1). The progress predicted by the model is $f(\hat{y}_k) - \bar{\omega}_{k+1}(y_{k+1}) = f(\hat{y}_k) - W_k(y_{k+1})$. Actual progress will be compared to a threshold value which arises from damping the progress predicted by the model by some $\kappa \in (0, 1)$

$$\vartheta_{k+1} = \kappa[f(\hat{y}_k) - \bar{\omega}_{k+1}(y_{k+1})].$$

Next, f is evaluated at y_{k+1} by calling the oracle which returns $f(y_{k+1})$ and a new minorant ω_{k+1} with $\omega_{k+1}(y_{k+1}) = f(y_{k+1})$. If progress in objective value is sufficiently large in comparison to the progress predicted by the model, i. e.,

$$f(\hat{y}_k) - f(y_{k+1}) \geq \vartheta_{k+1},$$

the method executes a *descent step* which moves the center of stability to the new point, $\hat{y}_{k+1} = y_{k+1}$. Otherwise, in a *null step*, the center remains unchanged, $\hat{y}_{k+1} = \hat{y}_k$, but the new minorant ω_{k+1} is used to improve the model. In fact, the requirement $\{\bar{\omega}_{k+1}, \omega_{k+1}\} \subseteq W_{k+1}$ ensures convergence of the function values $f(\hat{y}_k)$ to $f^* = \inf_{y \in C} f(y)$ under mild technical conditions on H_{k+1}. For these it suffices, *e.g.* to fix $0 < \underline{\lambda} \leq \bar{\lambda}$ and to choose $\underline{\lambda} I \preceq H_{k+1} \preceq \bar{\lambda} I$ following a descent step and $H_k \preceq H_{k+1} \preceq \bar{\lambda} I$ following a null step, see [5].

The decisive elements for an efficient implementation are the following:

- the choice of the cutting model W_k,
- the choice of the proximal term, in our case of H_k,
- the solution method for the bundle subproblem (1) with corresponding structural requirements on supported ground sets C.

and their interplay. While most bundle implementations employ polyhedral cutting models combined with a suitable active set QP approach, the ConicBundle callable library [15] is primarily designed for (nonnegative combinations of) conic cutting models built from symmetric cones. In particular the cone of positive semidefinite matrices and the second order cone lead to nonpolyhedral models that change significantly in each step. In solving (1) for these models, interior point methods are currently the best option available, but with these methods the cost of assembling the coefficients and solving the subproblem dominates the work per iteration for most applications. This paper explores possibilities to replace the classical Schur complement approach for computing the Newton step by an iterative approach in order to improve applicability to large scale problems.

In order to describe the main structure of the primal dual KKT-system for (1), let us briefly sketch the conic cutting models employed. These build on combinations of the cone of nonnegative vectors, the second order cone and the cone of positive semidefinite matrices, each with a specific trace vector for measuring the "size" of its elements,

$$\begin{aligned} &x \in \mathbb{R}^n_+ := \{x \in \mathbb{R}^n \colon x_i \geq 0, i = 1, \ldots, n\} && \Leftrightarrow\colon x \geq 0, && \text{trace vector } \mathbb{1}_n = \begin{bmatrix} 1 \\ \vdots \\ 1 \end{bmatrix}, \\ &x \in \mathcal{Q}^n := \left\{ \begin{pmatrix} x_1 \\ \bar{x} \end{pmatrix} \in \mathbb{R}^n \colon x_1 \geq \|\bar{x}\| \right\} && \Leftrightarrow\colon x \trianglerighteq 0, && \text{trace vector } e_1 = \begin{bmatrix} 1 \\ 0 \\ \vdots \\ 0 \end{bmatrix}, \\ &X \in \mathbb{S}^n_+ := \{X \in \mathbb{S}^n \colon \lambda_{\min}(X) \geq 0\} && \Leftrightarrow\colon X \succeq 0, && \text{trace matrix } I_n. \end{aligned}$$

Cartesian products of these are described by a triple $t = (n, q, s)$ with $n \in \mathbb{N}_0$, $q \in \mathbb{N}^{n_q}$, $s \in \mathbb{N}^{n_s}$ for some $n_q \in \mathbb{N}_0$, $n_s \in \mathbb{N}_0$ specifying the cone

$$\mathcal{S}_+^t := \mathbb{R}_+^n \times \bigtimes_{i=1}^{n_q} \mathcal{Q}^{q_i} \times \bigtimes_{i=1}^{n_s} \mathbb{S}_+^{s_i}. \tag{2}$$

The cone $\mathcal{S}_+^t$ will be regarded as a full dimensional cone in $\mathbb{R}^t := \mathbb{R}^{n(t)}$ with $n(t) = n + \mathbb{1}^\top q + \sum_{i=1}^{n_s} \binom{s_i+1}{2}$. Whenever convenient, an $x \in \mathcal{S}_+^t$ or $\mathbb{R}^t$ will be split into

$$x = (\xi^\top, (x^{(1)})^\top, \ldots, (x^{(n_q)})^\top, \operatorname{svec}(X^{(1)})^\top, \ldots, \operatorname{svec}(X^{(n_s)})^\top)^\top \tag{3}$$

in the natural way. Indices or elements may be omitted if the corresponding counters n, n_q, n_s happen to be 0 or 1. The *trace* $\operatorname{tr}(\cdot)$ of an element $x \in S_t$ is defined to be

$$\operatorname{tr} x := \langle \mathbb{1}_t, x \rangle := \langle \mathbb{1}, \xi \rangle + \sum_{i=1}^{n_q} \left\langle e_1, x^{(i)} \right\rangle + \sum_{i=1}^{n_s} \left\langle I, X^{(i)} \right\rangle.$$

In ConicBundle each cutting model may be considered to be specified by a tuple $M = (t, \alpha, K, \mathcal{B}, \underline{\omega})$ via

$$W_M(y) = \max_{\substack{x \in \mathcal{S}_+^t \\ \tau - \langle \mathbb{1}_t, x \rangle \in K}} [\underline{\omega}(y) + \langle \mathcal{B}(y), x \rangle]$$

where

$t = (n, q, s)$	specifies the cone as above,
$\tau > 0$	gives the trace value or trace upper bound,
$K \in \{\{0\}, \mathbb{R}_+\}$	specifies constant or bounded trace,
$\mathcal{B}\colon \mathbb{R}^m \to \mathbb{R}^{n(t)}$, $y \mapsto \mathcal{B}(y) = B_0 + By$	represents the bundle as affine function,
$\underline{\omega} = (\underline{\gamma}, \underline{g})$	provides a constant offset subgradient.

For example, the standard polyhedral model for h subgradients $\omega_i = (\gamma_i, g_i)$, $i = 1, \ldots, h$ is obtained for $M = (t = (h, 0, 0), \tau = 1, K = \{0\}, (B_0 = \begin{bmatrix} \gamma_1 \\ \vdots \\ \gamma_h \end{bmatrix}, B = \begin{bmatrix} g_1^\top \\ \vdots \\ g_h^\top \end{bmatrix}), \underline{\omega} = 0)$. Indeed, maximizing $\langle \mathcal{B}(y), x \rangle$ over $\xi \in \mathbb{R}_+^h$ with $\mathbb{1}^\top \xi = 1$ finds the best convex combination of the subgradients at y. In polyhedral models $\mathcal{B}$ may well be sparse, but in combination with positive semidefinite models large parts of $\mathcal{B}$ will be dense. Therefore we will not assume any specific structure in the bundle $\mathcal{B}$.

Depending on the variable metric heuristic in use, see [17, 18] for typical choices in ConicBundle, the proximal term may either be a positive multiple of I_m or it may be of the form $H = D_H + V_H V_H^\top \in \mathbb{S}_+^n$, where D_H is a diagonal matrix with strictly positive diagonal entries and $V_H \in \mathbb{R}^{m \times h_H}$ specifies a rank h_H contribution. If V_H is present, it typically consists of dense orthogonal columns.

The final ingredient is the basic optimization set C, which may have a polyhedral description of the form

$$C = \{y\colon \underline{a} \le Ay \le \overline{a}, \underline{y} \le y \le \overline{y}\},$$

where $A \in \mathbb{R}^{h_A \times m}$, $\underline{a} \in (\mathbb{R} \cup \{-\infty\})^{h_A}$, $\overline{a} \in (\mathbb{R} \cup \{\infty\})^{h_A}$, $\underline{y} \in (\mathbb{R} \cup \{-\infty\})^m$, $\overline{y} \in (\mathbb{R} \cup \{\infty\})^m$ are given data. If A is employed in applications, we expect the number of rows h_A of A to be small in comparison to m. The set C is tested for feasibility in advance, but no preprocessing is applied.

In order to reduce the indexing load in this presentation, we consider problem (1) in $\hat{y}$ only for a single iteration and one model $(t, \alpha, K, \mathcal{B}, \underline{\omega} = (\underline{\gamma}, \underline{g}))$. Putting $b = -H\hat{y} + \underline{g}$ and $\delta = \frac{1}{2}\langle H\hat{y}, \hat{y} \rangle + \underline{\gamma}$

the bundle problem may be written in the form

$$\begin{array}{ccc} \min & \max & \\ Ay - w = 0 & x \in \mathcal{S}^t_+ & \\ \underline{a} \le w \le \overline{a} & \alpha - \langle \mathbb{1}_t, x\rangle - \sigma = 0 & \\ \underline{y} \le y \le \overline{y} & \sigma \in K & \end{array} \qquad \left[\frac{1}{2}\langle Hy, y\rangle + \langle b, y\rangle + \langle B_0, x\rangle + \langle By, x\rangle + \delta\right]. \tag{4}$$

The existence of saddle points is guaranteed by compactness of the x-set and by strong convexity for y due to $H \succ 0$, see *e.g.* [22]. For the purpose of presentation, the primal dual KKT system for solving the saddle point problem (4) is built for $\underline{a} < \overline{a}$, $\underline{y} < \overline{y}$ and $K = \mathbb{R}_+$. The extension to the equality cases follows quite naturally and will be commented on at appropriate places. Throughout we will assume that $A, A^\top$ and $B, B^\top$ are given by matrix-vector multiplication oracles. A is assumed to have few rows, B may actually have a large number of columns, H is a positive diagonal plus low rank, but no further structural information is assumed to be available.

The original spectral bundle approach of [19] was designed for the unconstrained case $C = \mathbb{R}^m$ which allows direct elimination of y by convex optimality. Setting up the maximization problem for x then requires forming the typically dense Schur complement $BH^{-1}B^\top$. For increasing bundle sizes this is in fact the limiting operation within each bundle iteration. The aim of developing an iterative approach for (4) is therefore not only to allow for general C but also to circumvent the explicit computation of this Schur complement.

For setting up a primal-dual interior point approach for solving (4), the dual variables to constraints on the minimizing side will be denoted by $s \in \mathbb{R}^{h_A}$, $s_{\underline{a}}, s_{\overline{a}} \in \mathbb{R}^{h_A}_+$, $s_{\underline{y}}, s_{\overline{y}} \in \mathbb{R}^m_+$, the dual variables to the constraints on the maximizing side will be $z \in \mathcal{S}^t_+$, $z_\alpha \in \mathbb{R}$ and $\zeta \in K^* = \mathbb{R}_+$.

With barrier parameter $\mu > 0$ for the conic constraints the usual primal-dual KKT-System may be arranged in the form

$$\begin{array}{rcrcrcrcrcr} Hy & + & A^\top s & + & B^\top x & & & + & s_{\overline{y}} - s_{\underline{y}} & = & -b \\ Ay & & & & & & & - & w & = & 0 \\ By & & & & & - & \mathbb{1}_t\zeta & + & z & = & -B_0 \\ & & & - & \langle \mathbb{1}_t, x\rangle & & & - & \sigma & = & -\alpha \\ & - & s & & & & & + & s_{\overline{a}} - s_{\underline{a}} & = & 0 \end{array} \tag{5}$$

$$\begin{array}{rcl@{\qquad}rcl} (w - \underline{a}) \circ s_{\underline{a}} & = & \mu\mathbb{1} & (\overline{a} - w) \circ s_{\overline{a}} & = & \mu\mathbb{1} \\ (y - \underline{y}) \circ s_{\underline{y}} & = & \mu\mathbb{1} & (\overline{y} - y) \circ s_{\overline{y}} & = & \mu\mathbb{1} \\ x \circ_t z & = & \mu\mathbb{1}_t & \sigma\zeta & = & \mu. \end{array}$$

In this, "$\circ$" denotes the componentwise Hadamard product and "$\circ_t$" a canonical generalization to the cone $\mathcal{S}^t_+$, employing the arrow operator for second order cone parts and (typically symmetrized) matrix products for semidefinite parts.

In solving this by Newton's method, the linearization of the first perturbed complementarity line yields

$$\begin{array}{rcl} \Delta s_{\underline{a}} & = & \mu(w - \underline{a})^{-1} - s_{\underline{a}} - \Delta w \circ s_{\underline{a}} \circ (w - \underline{a})^{-1} \\ \Delta s_{\overline{a}} & = & \mu(\overline{a} - w)^{-1} - s_{\overline{a}} + \Delta w \circ s_{\overline{a}} \circ (\overline{a} - w)^{-1} \\ \hline \Delta s_{\overline{a}} - \Delta s_{\underline{a}} & = & \Delta w \circ \left((s_{\underline{a}} \circ (w - \underline{a})^{-1}) + (s_{\overline{a}} \circ (\overline{a} - w)^{-1})\right) \\ & & + s_{\underline{a}} - s_{\overline{a}} + \mu\left((\overline{a} - w)^{-1} - (w - \underline{a})^{-1}\right) \\ & =: & \Delta w \circ d_w + s_{\underline{a}} - s_{\overline{a}} - \mu c_w. \end{array}$$

Likewise, the second perturbed complementarity line gives rise to

$$\begin{array}{rcl} \Delta s_{\overline{b}} - \Delta s_{\underline{b}} & = & \Delta y \circ \left((s_{\underline{y}} \circ (y - \underline{y})^{-1}) + (s_{\overline{y}} \circ (\overline{y} - y)^{-1})\right) \\ & & + s_{\underline{y}} - s_{\overline{y}} + \mu\left((\overline{y} - y)^{-1} - (y - \underline{y})^{-1}\right) \\ & =: & \Delta y \circ d_y + s_{\underline{y}} - s_{\overline{y}} - \mu c_y. \end{array}$$

For dealing with the conic complementarity "$\circ_t$" we employ the symmetrization operators $\mathcal{E}_t$ and $\mathcal{F}_t$ of [30] in diagonal block form corresponding to $\mathcal{S}^t_+$, which give rise to a symmetric positive

definite $\mathfrak{X}_t = \mathcal{E}_t^{-1}\mathcal{F}_t \succ 0$ with diagonal block structure according to $\mathcal{S}_{++}^t$. With this the last perturbed complementarity line results in

$$\begin{aligned} \Delta z &= -\mathfrak{X}_t^{-1}\Delta x + \mu x^{-1} - z, \\ \Delta\sigma &= -\zeta^{-1}\sigma\Delta\zeta + \mu\zeta^{-1} - \sigma. \end{aligned}$$

Employing the linearization of the defining equation for s,

$$\Delta s_{\overline{a}} - \Delta s_{\underline{a}} = \Delta s + s + s_{\underline{a}} - s_{\overline{a}},$$

the variable Δw may now be eliminated via

$$\Delta w = d_w^{-1} \circ \Delta s + d_w^{-1} \circ s + \mu d_w^{-1} \circ c_w.$$

Put $D_y := \mathrm{Diag}(d_y) > 0$ and $D_w := \mathrm{Diag}(d_w) > 0$, then the Newton step is obtained by solving the system

$$\begin{bmatrix} H + D_y & A^\top & B^\top & 0 \\ A & -D_w^{-1} & 0 & 0 \\ B & 0 & -\mathfrak{X}_t^{-1} & -\mathbb{1}_t \\ 0 & 0 & -\mathbb{1}_t^\top & \zeta^{-1}\sigma \end{bmatrix} \begin{pmatrix} \Delta y \\ \Delta s \\ \Delta x \\ \Delta\zeta \end{pmatrix} = \begin{pmatrix} r_y \\ r_s \\ r_x \\ r_\zeta \end{pmatrix}, \tag{6}$$

with right hand side

$$\begin{pmatrix} r_y \\ r_s \\ r_x \\ r_\zeta \end{pmatrix} = \begin{pmatrix} -(Hy + B^\top x + A^\top s + b) & + & \mu c_y \\ -(Ay - w) + d_w^{-1} \circ s & + & \mu d_w^{-1} \circ c_w \\ -(By + B_0 - \mathbb{1}_t\zeta) & - & \mu x^{-1} \\ -(\alpha - \langle \mathbb{1}_t, x\rangle) & + & \mu\zeta^{-1} \end{pmatrix}.$$

The right hand side may be modified as usual to obtain predictor and corrector right hand sides, but this will not be elaborated on here. Note, for $K = \{0\}$ the same system works with $\sigma = 0$ and without the centering terms associated with ζ in r_ζ. Likewise, whenever line i of A corresponds to an equation, i. e., $\underline{a}_i = \overline{a}_i$, the respective entry of d_w^{-1} has to be replaced by zero.

Eq. (6) is a symmetric indefinite system, that could be solved by appropriate iterative methods directly. So far, however, we were not able to conceive suitable general preconditioning approaches for exploiting the given structural properties in the full system. Surprisingly, a viable path seems to be offered by the traditional Schur complement approach after all. The resulting system allows to perform matrix vector multiplications at minimal additional cost and is frequently positive definite.

To see this, first take the Schur complement with respect to the $\mathfrak{X}_t^{-1}$ block,

$$\begin{bmatrix} H + D_y + B^\top\mathfrak{X}_t B & A^\top & -B^\top\mathfrak{X}_t\mathbb{1}_t \\ A & -D_w^{-1} & 0 \\ -(B^\top\mathfrak{X}_t\mathbb{1}_t)^\top & 0 & \zeta^{-1}\sigma + \mathbb{1}_t^\top\mathfrak{X}_t\mathbb{1}_t \end{bmatrix}.$$

Assuming $d_w > 0$ (no equality constraint rows in A), eliminate the second and third block with further Schur complements and split $B^\top\mathfrak{X}_t B = B^\top\mathfrak{X}_t^{\frac{1}{2}}\mathfrak{X}_t^{\frac{1}{2}}B$,

$$\tilde{H} = \left[H + D_y + B^\top\mathfrak{X}_t^{\frac{1}{2}}(I - \tfrac{\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t(\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top\mathfrak{X}_t\mathbb{1}_t})\mathfrak{X}_t^{\frac{1}{2}}B + A^\top D_w A\right]. \tag{7}$$

Also in the equality case of $\sigma = 0$ the matrix $I - \frac{\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t(\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top\mathfrak{X}_t\mathbb{1}_t} \succeq 0$ is positive semidefinite, so the resulting system is positive definite. Equality constraints in A induce zero diagonal elements in D_w (or in $-D_w^{-1}$). In this case the corresponding rows will not be eliminated and give rise to an indefinite system of the form $\left[\begin{smallmatrix} \tilde{H} & \tilde{A}^\top \\ \tilde{A} & 0 \end{smallmatrix}\right]$ with a large positive definite block $\tilde{H}$ and hopefully few further rows in $\tilde{A}$. For such systems it is well studied how to employ a preconditioner for $\tilde{H}$ to solve the full indefinite system with *e.g.* MINRES, see [10].

The cost of multiplying the full KKT-matrix of (6) by a vector is roughly the same as that of multiplying $\tilde{H}$ by a vector. Indeed, the same multiplications arise for $H + D_y, A, A^\top, B, B^\top$. So it remains to compare the cost of a multiplication by $\begin{bmatrix} -\mathfrak{X}^{-1} & -\mathbb{1}_t \\ -\mathbb{1}_t^\top & \zeta^{-1}\sigma \end{bmatrix}$ to a multiplication by $\mathfrak{X}_t^{\frac{1}{2}}(I - \frac{\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t(\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t)^\top}{\zeta^{-1}\sigma+\mathbb{1}_t^\top\mathfrak{X}_t\mathbb{1}_t})\mathfrak{X}_t^{\frac{1}{2}} = \mathfrak{X}_t - \frac{\mathfrak{X}_t\mathbb{1}_t(\mathfrak{X}_t\mathbb{1}_t)^\top}{\zeta^{-1}\sigma+\mathbb{1}_t^\top\mathfrak{X}_t\mathbb{1}_t}$. Recall that $\mathfrak{X}_t$ is a block diagonal matrix with a separate block for each cone $\mathbb{R}_+$, $\mathcal{Q}^{q_i}$, $\mathbb{S}_+^{s_i}$ specified by t and the cost of multiplying by $\mathfrak{X}_t$ or $\mathfrak{X}_t^{-1}$ is identical. Thus the only difference are the multiplications by $\mathbb{1}_t$ in the first case and by the precomputed vector $\mathfrak{X}_t\mathbb{1}_t$ in the second. The vector $\mathfrak{X}_t\mathbb{1}_t$ may be formed at almost negligible cost along with setting up $\mathfrak{X}_t$. So there is no noteworthy difference in the cost of matrix vector multiplications between the two systems and no structural advantages are lost when working with $\tilde{H}$ instead of the full system. We will therefore concentrate on developing a preconditioner for $\tilde{H}$.

For this note that $\tilde{H}$ of (7) arises from adding a Gram matrix to a positive diagonal,

$$\tilde{H} = D + VV^\top,$$

where (recall $H = D_H + V_H V_H^\top$)

$$D = D_H + D_y \quad \text{and} \quad V = \left[V_H, A^\top D_w^{\frac{1}{2}}, B^\top \mathfrak{X}_t^{\frac{1}{2}}(I - \frac{\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t(\mathfrak{X}_t^{\frac{1}{2}}\mathbb{1}_t)^\top}{\zeta^{-1}\sigma+\mathbb{1}_t^\top\mathfrak{X}_t\mathbb{1}_t})^{\frac{1}{2}}\right]. \tag{8}$$

Note that the multiplication of V with a vector requires only a little bit more than half the number of operations of multiplying $\tilde{H}$ (or the full KKT matrix) with a vector. This suggests to explore possibilities of finding low rank approximations of V for preconditioning.

3 Low rank preconditioning a Gram-matrix plus positive diagonal

Consider a matrix

$$H = D + VV^\top$$

with a positive definite matrix $D \in \mathbb{S}_{++}^m$ and $V \in \mathbb{R}^{m\times n}$. In our application D is diagonal, but the results apply for general $D \succ 0$. This is applicable in practice as long as D^{-1} can be applied efficiently to vectors. Matrix V is assumed to be given by a matrix-vector multiplication oracle, i. e., V and $V^\top$ may be multiplied by vectors but the matrix does not have to be available explicitly.

For motivating the following preconditioning approaches, first consider computing the singular value decomposition of

$$D^{-\frac{1}{2}}V = Q_H \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} P_H^\top$$

with orthogonal $Q_H \in \mathbb{R}^{m\times m}$, diagonal $\Sigma = \mathrm{Diag}(\sigma_1, \ldots, \sigma_n)$ ordered nonincreasingly by $\sigma_1 \geq \cdots \geq \sigma_n \geq 0$ and orthogonal $P_H \in \mathbb{R}^{n\times n}$ (for convenience, it is assumed that $n \leq m$). Then

$$H = D^{\frac{1}{2}}Q_H \begin{bmatrix} I_n + \Sigma^2 & 0 \\ 0 & I_{m-n} \end{bmatrix} Q_H^\top D^{\frac{1}{2}},$$

$$H^{-1} = D^{-\frac{1}{2}}Q_H \begin{bmatrix} (I_n + \Sigma^2)^{-1} & 0 \\ 0 & I_{m-n} \end{bmatrix} Q_H^\top D^{-\frac{1}{2}}.$$

When Σ is replaced by the k largest singular values, this gives rise to a good "low rank" preconditioner, see Theorem 1 below. Computing the full matrix $D^{-\frac{1}{2}}V$ and its singular value decomposition will in general be too costly or even impossible. Instead the general idea is to work with $D^{-\frac{1}{2}}V\Omega$ for some random or deterministic choice of $\Omega \in \mathbb{R}^{n\times k}$.

Multiplying by a random Ω may be thought of as giving rise to a subspace approximation in the style of Johnson-Lindenstrauss, cf. [1, 8], and this formed the starting point of this investigation. The actual randomized approach and analysis, however, mainly builds on [13] and the bounding

techniques presented there. For the deterministic preconditioning variant the recent work [12] provided strong guidance for analysing the condition number.

Here, Ω will mostly consist of orthonormal columns. Yet it is instructive to consider more general cases, as well. An arbitrary $\Omega \in \mathbb{R}^{n\times k}$ gives rise to the preconditioner

$$H_\Omega = D + V\Omega\Omega^\top V = D^{\frac{1}{2}} Q_H \begin{bmatrix} I_n + \Sigma P_H^\top \Omega\Omega^\top P_H \Sigma & 0 \\ 0 & I_{m-n} \end{bmatrix} Q_H^\top D^{\frac{1}{2}}. \tag{9}$$

Putting $G_\Omega = D^{\frac{1}{2}} Q_H \begin{bmatrix} I_n+\Sigma P_H^\top \Omega\Omega^\top P_H\Sigma & 0 \\ 0 & I_{m-n}\end{bmatrix}^{\frac{1}{2}}$ we have $H_\Omega = G_\Omega G_\Omega^\top$. The preconditioner is the better the closer $G_\Omega^{-1} H G_\Omega^{-T}$ is to the identity. In the analysis of convergence rates, see *e.g.* [10], this enters via the condition number

$$\kappa_\Omega := \frac{\lambda_{\max}(G_\Omega^{-1} H G_\Omega^{-T})}{\lambda_{\min}(G_\Omega^{-1} H G_\Omega^{-T})} = \frac{\lambda_{\max}(H^{\frac{1}{2}} H_\Omega^{-1} H^{\frac{1}{2}})}{\lambda_{\min}(H^{\frac{1}{2}} H_\Omega^{-1} H^{\frac{1}{2}})} = \frac{\lambda_{\max}(H_\Omega^{-\frac{1}{2}} H H_\Omega^{-\frac{1}{2}})}{\lambda_{\min}(H_\Omega^{-\frac{1}{2}} H H_\Omega^{-\frac{1}{2}})}.$$

In this, the equations follow from $BB^\top$ and $B^\top B$ having the same eigenvalues for $B \in \mathbb{R}^{n\times n}$.

Theorem 1 *Let $H = D + VV^\top \in \mathbb{S}^m_{++}$ with positive definite $D \in \mathbb{S}^m_{++}$ and $V \in \mathbb{R}^{m\times n}$ with $n < m$ and singular value decomposition $D^{-\frac{1}{2}} V = Q_H \Sigma P_H^\top$, $Q_H^\top Q_H = I_m$, $P_H^\top P_H = I_n$, $\Sigma = \mathrm{Diag}(\sigma_1 \geq \cdots \geq \sigma_n) \in \mathbb{S}^n_+$. For $\Omega \in \mathbb{R}^{n\times k}$ the preconditioner H_Ω of (9) results in condition number*

$$\begin{aligned} \kappa_\Omega &= \frac{\max\{1, \lambda_{\max}((I_n+\Sigma^2)^{\frac{1}{2}}(I_n + \Sigma P_H^\top \Omega\Omega^\top P_H \Sigma)^{-1}(I_n+\Sigma^2)^{\frac{1}{2}})\}}{\min\{1, \lambda_{\min}((I_n+\Sigma^2)^{\frac{1}{2}}(I_n + \Sigma P_H^\top \Omega\Omega^\top P_H \Sigma)^{-1}(I_n+\Sigma^2)^{\frac{1}{2}})\}} \\ &= \frac{\max\{1, \lambda_{\max}((I_n+\Sigma^2)^{-\frac{1}{2}}(I_n + \Sigma P_H^\top \Omega\Omega^\top P_H \Sigma)(I_n+\Sigma^2)^{-\frac{1}{2}})\}}{\min\{1, \lambda_{\min}((I_n+\Sigma^2)^{-\frac{1}{2}}(I_n + \Sigma P_H^\top \Omega\Omega^\top P_H \Sigma)(I_n+\Sigma^2)^{-\frac{1}{2}})\}} \end{aligned}$$

In particular, for $0 \leq k < n$ and $\Omega = (P_H)_{\bullet,[1,\ldots,k]}$ the condition number's value is $1+\sigma_{k+1}^2$.

Proof. For G_Ω as above direct computation yields

$$\begin{aligned} G_\Omega^{-1} H G_\Omega^{-T} &= \begin{bmatrix} I_n+\Sigma\Omega\Omega^\top\Sigma & 0 \\ 0 & I_{m-n}\end{bmatrix}^{-\frac{1}{2}} \begin{bmatrix} I_n+\Sigma^2 & 0 \\ 0 & I_{m-n}\end{bmatrix} \begin{bmatrix} I_n+\Sigma\Omega\Omega^\top\Sigma & 0 \\ 0 & I_{m-n}\end{bmatrix}^{-\frac{1}{2}} \\ &= \begin{bmatrix} (I_n+\Sigma\Omega\Omega^\top\Sigma)^{-\frac{1}{2}}(I_n+\Sigma^2)(I_n+\Sigma\Omega\Omega^\top\Sigma)^{-\frac{1}{2}} & 0 \\ 0 & I_{m-n}\end{bmatrix}. \end{aligned}$$

The eigenvalues of $(I_n + \Sigma\Omega\Omega^\top\Sigma)^{-\frac{1}{2}}(I_n+\Sigma^2)(I_n+\Sigma\Omega\Omega^\top\Sigma)^{-\frac{1}{2}}$ coincide with those of $(I_n + \Sigma^2)^{\frac{1}{2}}(I_n+\Sigma\Omega\Omega^\top\Sigma)^{-1}(I_n+\Sigma^2)^{\frac{1}{2}}$, because for $B = (I_n+\Sigma\Omega\Omega^\top\Sigma)^{-\frac{1}{2}}(I_n+\Sigma^2)^{\frac{1}{2}}$ the two matrices are $BB^\top$ and $B^\top B$. This gives rise to the first line. The second follows because for positive definite A there holds $\lambda_{\max}(A) = 1/\lambda_{\min}(A^{-1})$ and $\lambda_{\min}(A) = 1/\lambda_{\max}(A^{-1})$. ■

Consider now a subspace spanned by k orthonormal columns collected in some matrix $P_\ell \in \mathbb{R}^{n\times k}$ which hopefully generates most of the large directions of $D^{-\frac{1}{2}}V$. In this orthonormal case a simpler bound on the condition number may be obtained by following the argument of Th. 5.1 in [12].

Theorem 2 *Let $H = D + VV^\top$ with positive definite $D \in \mathbb{S}^m_{++}$ and general $V \in \mathbb{R}^{m\times n}$, let $P = [P_\ell, P_s] \in \mathbb{R}^{n\times n}$, $PP^\top = I_n$. Preconditioner $H_{P_\ell} = D + VP_\ell P_\ell^\top V^\top$ has condition number $\kappa_{P_\ell} \leq 1 + \lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} V P_s P_s^\top V^\top H_{P_\ell}^{-\frac{1}{2}})$. Equality holds if and only if $\mathrm{rank}(VP_s) < m$.*

Proof. Because $H = D + [VP_\ell, VP_s][VP_\ell, VP_s]^\top = H_{P_\ell} + VP_sP_s^\top V^\top$ we have

$$H_{P_\ell}^{-\frac{1}{2}} H H_{P_\ell}^{-\frac{1}{2}} = I_m + H_{P_\ell}^{-\frac{1}{2}} V P_s P_s^\top V^\top H_{P_\ell}^{-\frac{1}{2}}.$$

The second summand is positive semidefinite with minimum eigenvalue 0 if and only if $\mathrm{rank}(VP_s) < m$. Thus,

$$\begin{aligned} \lambda_{\min}(H_{P_\ell}^{-\frac{1}{2}} H H_{P_\ell}^{-\frac{1}{2}}) &\geq 1, \\ \lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} H H_{P_\ell}^{-\frac{1}{2}}) &= 1 + \lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} V P_s P_s^\top V^\top H_{P_\ell}^{-\frac{1}{2}}). \end{aligned}$$

By $\kappa(H_{P_\ell}^{-\frac{1}{2}} H H_{P_\ell}^{-\frac{1}{2}}) = \frac{\lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} H H_{P_\ell}^{-\frac{1}{2}})}{\lambda_{\min}(H_{P_\ell}^{-\frac{1}{2}} H H_{P_\ell})}$ the result is proved. ■

Building on these two theorems we first analyze randomized approaches that do not make any assumptions on structural properties of $D^{-\frac{1}{2}}V$ but only require a multiplication oracle. Afterwards we present a deterministic approach that exploits some knowledge of the bundle subproblem and the interior point algorithm. The corresponding routines supply a $\hat{V} = V\Omega$. The actual preconditioning routine, Alg. 3 below, does not use H_Ω directly, but a truncated preconditioner $H_{\Omega\hat{P}}$ that drops all singular values of $D^{-\frac{1}{2}}V\Omega$ that are less than one. The inverse is then formed via a Woodbury-formula, see [23, §0.7.4]. Note, depending on the expected number of calls to the routine and the structure preserved in $\hat{V}$, it may or may not pay off to also precompute $\hat{V}\hat{P}$. For diagonal D and dense $\hat{V}$ the cost of applying this preconditioner is $O(m + mk + k\hat{k})$.

Algorithm 3 (Preconditioning by truncated $H_\Omega = D + V\Omega(V\Omega)^\top$)
Input: *$v \in \mathbb{R}^m$, $D \in \mathbb{S}^n_{++}$, precomputed $\hat{V} = V\Omega \in \mathbb{R}^{n\times k}$ and, for $\hat{V}^\top D^{-1}\hat{V} = P\,\mathrm{Diag}(\hat{\lambda}_1 \geq \cdots \geq \hat{\lambda}_k)P^\top$, $\hat{k} = \max\{0, i\colon \hat{\lambda}_i \geq 1\}$, $\hat{\Lambda} = \mathrm{Diag}(\hat{\lambda}_1, \ldots, \hat{\lambda}_{\hat{k}})$, $\hat{P} = P_{\bullet,[1,\ldots,\hat{k}]}$.*
Output: *$(H_{\Omega\hat{P}})^{-1}v$.*
1. $v \leftarrow D^{-1}v$.
2. If $\hat{k} > 0$ set $v \leftarrow v - D^{-1}\hat{V}\hat{P}(I + \hat{\Lambda})^{-1}\hat{P}^\top\hat{V}^\top v$.
3. return v.

3.1 Preconditioning by Random Subspaces

For the random subspace approach fix some $k \in \mathbb{N}$ with $2 \leq k < n$, let Ω be a $n \times k$ random matrix whose elements are independently identically distributed by the normal distribution $\mathcal{N}(0, \frac{1}{k})$. For this Ω consider first the low rank approximation $D^{-\frac{1}{2}}V\Omega = Q_H \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} P_H^\top \Omega$. Because the normal distribution is invariant under orthogonal transformations, we may assume $P_H = I$ and analyze the setting $Q_H \begin{bmatrix} \Sigma \\ 0 \end{bmatrix} \Omega$ giving rise to the low rank approximation by the random matrix

$$H_\Omega = D^{\frac{1}{2}} Q_H \begin{bmatrix} I_n + \Sigma\Omega\Omega^\top\Sigma & 0 \\ 0 & I_{m-n} \end{bmatrix} Q_H^\top D^{\frac{1}{2}}.$$

In view of Theorem 1 such a preconditioner is good if $(I + \Sigma^2)^{-\frac{1}{2}}(I + \Sigma\Omega\Omega^\top\Sigma)(I + \Sigma^2)^{-\frac{1}{2}}$ is close to the identity. Based on the Johnson-Lindenstrauss interpretation, it seems likely that large portions of the spectrum will be close to one. This can be justified to some extent by studying the moments of the Ritz values.

Theorem 4 *Let $\Omega \in \mathbb{R}^{n\times k}$ have its elements i.i.d. according to the normal distribution $\mathcal{N}(0, \frac{1}{k})$, then for any $x \in \mathbb{R}^n$ the quadratic form*

$$q(x) = x^\top(I + \Sigma^2)^{-\frac{1}{2}}(I + \Sigma\Omega\Omega^\top\Sigma)(I + \Sigma^2)^{-\frac{1}{2}}x$$

has expected value $\mathbb{E}(q(x)) = \|x\|^2$ and variance $\mathrm{Var}(q(x)) = \frac{2}{k}\big(\sum_{i=1}^n \frac{\sigma_i^2}{1+\sigma_i^2}x_i^2\big)^2$.

Proof. Let $\Omega = (\omega_{ij})$ with i.i.d. elements ω_{ij} from $\mathcal{N}(0, \frac{1}{k})$. Recall that $\mathbb{E}(\omega_{ij}) = 0$, $\mathbb{E}(\omega_{ij}^2) = \frac{1}{k}$, $\mathbb{E}(\omega_{ij}^3) = 0$, $\mathbb{E}(\omega_{ij}^4) = 3/k^2$ and that for independent random variables X, Y there holds $\mathbb{E}(XY) = \mathbb{E}(X)\mathbb{E}(Y)$.

The expected value of the quadratic form evaluates to

$$\begin{aligned}
&\mathbb{E}\Big(x^\top (I+\Sigma^2)^{-\frac{1}{2}}(I+\Sigma\Omega\Omega^\top\Sigma)(I+\Sigma^2)^{-\frac{1}{2}}x\Big) = \\
&= \sum_{i=1}^n \frac{1}{1+\sigma_i^2}x_i^2 + \mathbb{E}\Big(\sum_{h=1}^k \big(\sum_{i=1}^n \frac{\sigma_i}{\sqrt{1+\sigma_i^2}}\omega_{hi}x_i\big)^2\Big) \\
&= \sum_{i=1}^n \frac{1}{1+\sigma_i^2}x_i^2 + \sum_{h=1}^k\sum_{i=1}^n \frac{\sigma_i^2}{1+\sigma_i^2}x_i^2\mathbb{E}\omega_{hi}^2 \\
&= \sum_{i=1}^n x_i^2\Big(\frac{1}{1+\sigma_i^2} + \frac{\sigma_i^2}{1+\sigma_i^2}\sum_{h=1}^k \mathbb{E}\omega_{hi}^2\Big) \\
&= \sum_{i=1}^n x_i^2\Big(\frac{1}{1+\sigma_i^2} + \frac{\sigma_i^2}{1+\sigma_i^2}\sum_{h=1}^k \frac{1}{k}\Big) = \sum_{i=1}^n x_i^2 = \|x\|^2.
\end{aligned}$$

For determining the variance, the second moment may be computed as follows.

$$\begin{aligned}
&\mathbb{E}\Big(\Big[\sum_{h=1}^k \big(\sum_{i=1}^n \frac{\sigma_i}{\sqrt{1+\sigma_i^2}}\omega_{hi}x_i\big)^2\Big]^2\Big) = \\
&= \mathbb{E}\Big(\Big[\sum_{h=1}^k\sum_{i=1}^n\sum_{i'=1}^n \frac{\sigma_i}{\sqrt{1+\sigma_i^2}}\frac{\sigma_{i'}}{\sqrt{1+\sigma_{i'}^2}}\omega_{hi}x_i\omega_{hi'}x_{i'}\Big]^2\Big) \\
&= \mathbb{E}\Big(\sum_{h,h'=1}^k \sum_{i,i',j,j'=1}^n \\
&\qquad \frac{\sigma_i}{\sqrt{1+\sigma_i^2}}\frac{\sigma_{i'}}{\sqrt{1+\sigma_{i'}^2}}\frac{\sigma_j}{\sqrt{1+\sigma_j^2}}\frac{\sigma_{j'}}{\sqrt{1+\sigma_{j'}^2}}x_ix_{i'}x_jx_{j'}\omega_{hi}\omega_{hi'}\omega_{h'j}\omega_{h'j'}\Big).
\end{aligned}$$

In the cases of $h \neq h'$ only terms with $i = i'$ and $j = j'$ are not zero. These evaluate to $\mathbb{E}\omega_{hi}^2\mathbb{E}\omega_{h'j}^2 = \frac{1}{k^2}$ giving

$$\frac{k(k-1)}{k^2}\sum_{i=1}^n\sum_{j=1}^n \frac{\sigma_i^2}{1+\sigma_i^2}x_i^2\frac{\sigma_j^2}{1+\sigma_j^2}x_j^2.$$

For each $h = h'$ there remain ($i = i' = j = j'$) with value $\mathbb{E}\omega_{hi}^4 = \frac{3}{k^2}$,

$$\frac{3k}{k^2}\sum_{i=1}^n \Big(\frac{\sigma_i^2}{1+\sigma_i^2}\Big)^2 x_i^4,$$

and the three pairings $(i = i', j = j')$, $(i = j, i' = j')$ and $(i = j', i' = j)$ each with value $\frac{1}{k^2}$,

$$\frac{3k}{k^2}\sum_{i\neq j} \frac{\sigma_i^2}{1+\sigma_i^2}x_i^2\frac{\sigma_j^2}{1+\sigma_j^2}x_j^2.$$

Summing up these three expressions yields

$$\mathbb{E}\Big(\Big[\sum_{h=1}^k \big(\sum_{i=1}^n \frac{\sigma_i}{\sqrt{1+\sigma_i^2}}\omega_{hi}x_i\big)^2\Big]^2\Big) = \frac{k^2+2k}{k^2}\Big(\sum_{i=1}^n \frac{\sigma_i^2}{1+\sigma_i^2}x_i^2\Big)^2.$$

The result now follows from the usual $\operatorname{Var} X = \mathbb{E}(X^2) - (\mathbb{E}X)^2$ for any random variable X. ■

This suggests that even for relatively small k the behavior of the preconditioned system may be expected to be reasonably close to the identity for a large portion of the directions. The result, however, does not seem to open a path towards good bounds on the condition number.

A first possibility is offered by Theorem 2. Recall that for an arbitrary matrix $A \in \mathbb{R}^{m\times n}$ the projector

$$\mathbf{P}_A = A(A^\top A)^\dagger A^\top$$

projects any vector of $\mathbb{R}^m$ onto the range space of A and $\mathbf{P}$ depends only on this range space. Computationally it may be determined by a QR-decomposition of $A = Q_A R_A$ with orthogonal $Q_A \in \mathbb{R}^{m\times n'}$ for $n' = \operatorname{rank}(A)$ via $\mathbf{P}_A = Q_A Q_A^\top$. The formula allows to verify $\mathbf{P}_A = \mathbf{P}_A^\top$, $\mathbf{P}_A\mathbf{P}_A = \mathbf{P}_A$ and $\mathbf{P}_A A = A$ by direct computation. Furthermore, for any $B \in \mathbb{R}^{n\times h}$ there holds $\mathbf{P}_{AB} \preceq \mathbf{P}_A \preceq I_m$ because of the containment relations between the ranges.

In the following $\Omega\Omega^\top$ in H_Ω will be replaced by the projector $\mathbf{P}_\Omega$. The random low rank approximation to be considered reads

$$H_{P_\Omega} = D^{\frac{1}{2}} Q_H \begin{bmatrix} I_n + \Sigma\mathbf{P}_\Omega\Sigma & 0 \\ 0 & I_{m-n} \end{bmatrix} Q_H^\top D^{\frac{1}{2}}.$$

Corollary 5 *Let $H = D + VV^\top \in \mathbb{R}^m$ with positive definite $D \in \mathbb{S}^m_{++}$ and $V \in \mathbb{R}^{m\times n}$. Given $\Omega \in \mathbb{R}^{n\times k}$, let $\mathbf{P}_\Omega = \Omega(\Omega^\top\Omega)^\dagger\Omega^\top$. For the preconditioner $H_{\mathbf{P}_\Omega}$ the condition number satisfies $\kappa_{\mathbf{P}_\Omega} \le 1 + \|D^{-\frac{1}{2}}V(I - \mathbf{P}_\Omega)\|^2$, where $\|\cdot\|$ denotes the spectral norm.*

Proof. Let $\mathbf{P}_\Omega = P_\ell P_\ell^\top$ with $P_\ell \in \mathbb{R}^{n\times k'}$ for some $k' \le k$ and $P_\ell^\top P_\ell = I_{k'}$. Add orthonormal columns P_s so that $P = [P_\ell, P_s]$ satisfies $PP^\top = I_n$. Note, $I - \mathbf{P}_\Omega = P_s P_s^\top$ is the projector onto the orthogonal complement. Use this choice in Theorem 2, then $H_{\mathbf{P}_\Omega} = H_{P_\ell} \succeq D$. Observe that by $D \in \mathbb{S}^m_{++}$ for any $\lambda \ge 0$, $A \in \mathbb{S}^m$ the relation $D^{-\frac{1}{2}}AD^{-\frac{1}{2}} \preceq \lambda I_m$ is equivalent to $A \preceq \lambda D$ and implies $A \preceq \lambda(D + VP_\ell P_\ell^\top V^\top) = \lambda H_{P\ell}$ which is equivalent to $H_{P_\ell}^{-\frac{1}{2}} A H_{P_\ell}^{-\frac{1}{2}} \preceq \lambda I_m$. Therefore

$$\lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} V P_s P_s^\top V^\top H_{P_\ell}^{-\frac{1}{2}}) \le \lambda_{\max}(D^{-\frac{1}{2}} V P_s P_s^\top V^\top D^{-\frac{1}{2}}) = \|D^{-\frac{1}{2}}V(I - \mathbf{P}_\Omega)\|^2.$$

■

While this bound is rather straight forward to derive, it does not seem strong enough to observe a reduced influence of the largest singular values of $D^{-\frac{1}{2}}V$. Indeed, in its derivation only the diagonal of $H_{\mathbf{P}_\Omega}$ was considered and the influence of $V\Omega$ is lost.

In order to obtain stronger bounds, the rather involved techniques layed out in [13] seem to be required. The next steps and results follow their arguments closely. This time the H_{P_Ω}-part is kept inverted in the analysis of the condition number.

Because $I + \Sigma P_\Omega \Sigma \preceq I + \Sigma^2$ there holds

$$(I + \Sigma^2)^{\frac{1}{2}}(I + \Sigma P_\Omega \Sigma)^{-1}(I + \Sigma^2)^{\frac{1}{2}} \succeq I.$$

By Theorem 1 the condition number is bounded by

$$\kappa_{P_\Omega} \le \lambda_{\max}((I + \Sigma^2)^{\frac{1}{2}}(I + \Sigma P_\Omega \Sigma)^{-1}(I + \Sigma^2)^{\frac{1}{2}})$$

and will attain it, whenever $n < m$. In terms of Ω, the best possible outcome is an event resulting in $P_\Omega = \begin{bmatrix} I_k \\ 0 \end{bmatrix}$ (see, *e.g.* [23, 7.4.52]). It corresponds to the truncated SVD and gives $\kappa_{\left[\begin{smallmatrix} I_k \\ 0 \end{smallmatrix}\right]} = 1 + \sigma_{k+1}^2$. Aiming for something more realistic, one hopes for a good coverage of the first k singular values when oversampling with p additional columns.

The first step in the analysis is to obtain a deterministic bound for a fixed $\Omega \in \mathbb{R}^{n\times(k+p)}$ as outlined in [13, §9.2].

Theorem 6 *Given $\sigma_1 \geq \cdots \geq \sigma_n \geq 0$ and a matrix $\Omega \in \mathbb{R}^{n\times(k+p)}$ with $k \leq n$ so that the first k rows of Ω are linearly independent, split $\Sigma = \begin{bmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{bmatrix}$ into blocks $\Sigma_1 = \mathrm{Diag}(\sigma_1, \ldots, \sigma_k)$ and $\Sigma_2 = \mathrm{Diag}(\sigma_{k+1}, \ldots, \sigma_n)$ and $\Omega = \begin{bmatrix} \Omega_1 \\ \Omega_2 \end{bmatrix}$ into the first k rows $\Omega_1 \in \mathbb{R}^{k\times(k+p)}$ and the last $n-k$ rows $\Omega_2 \in \mathbb{R}^{(n-k)\times k+p}$. Then*

$$\lambda_{\max}((I+\Sigma^2)^{\frac{1}{2}}(I+\Sigma P_\Omega \Sigma)^{-1}(I+\Sigma^2)^{\frac{1}{2}}) \leq 2 + \sigma_{k+1}^2 + \|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2.$$

Proof. By assumption Ω_1 has full row rank and the range space of the matrix

$$Z = \Omega \cdot \Omega_1^\dagger = \begin{bmatrix} I_k \\ F = \Omega_2\Omega_1^\dagger \end{bmatrix} \in \mathbb{R}^{n\times k}$$

is contained in the range space of Ω. Hence $P_Z \preceq P_\Omega$ and

$$\lambda_{\max}((I+\Sigma^2)^{\frac{1}{2}}(I+\Sigma P_\Omega \Sigma)^{-1}(I+\Sigma^2)^{\frac{1}{2}}) \leq \lambda_{\max}((I+\Sigma^2)^{\frac{1}{2}}(I+\Sigma P_Z \Sigma)^{-1}(I+\Sigma^2)^{\frac{1}{2}}). \tag{10}$$

The projector P_Z computes to

$$P_Z = \begin{bmatrix} I \\ F \end{bmatrix} \left[\, I + F^\top F \,\right]^{-1} \begin{bmatrix} I \\ F \end{bmatrix}^\top.$$

Use this in the Woodbury-formula for inverses of rank adjustments [23, §0.7.4] for $(I+\Sigma P_Z\Sigma)^{-1}$ to obtain

$$\begin{aligned}
(I+\Sigma P_Z\Sigma)^{-1} &= I - \begin{bmatrix} \Sigma_1 \\ \Sigma_2 F \end{bmatrix} \left[\, I + F^\top F + \Sigma_1^2 + F^\top \Sigma_2^2 F \,\right]^{-1} \begin{bmatrix} \Sigma_1 \\ \Sigma_2 F \end{bmatrix}^\top \\
&= I - \begin{bmatrix} \Sigma_1 \\ \Sigma_1 F \end{bmatrix} \left[\, I + \Sigma_1^2 + F^\top (I+\Sigma_2^2) F \,\right]^{-1} \begin{bmatrix} \Sigma_1 \\ \Sigma_2 F \end{bmatrix}^\top \\
&\preceq I - \begin{bmatrix} \Sigma_1 \\ \Sigma_2 F \end{bmatrix} \left[\, I + \Sigma_1^2 + \|(I+\Sigma_2^2)^{\frac{1}{2}} F\|^2 I \,\right]^{-1} \begin{bmatrix} \Sigma_1 \\ \Sigma_2 F \end{bmatrix}^\top.
\end{aligned} \tag{11}$$

The last line follows, because $\|(I+\Sigma_2^2)^{\frac{1}{2}}F\|^2 = \lambda_{\max}(F^\top(I+\Sigma_2^2)F) =: \bar{\lambda}$ and therefore $F^\top(I+\Sigma_2^2)F \preceq \bar{\lambda} I$ giving

$$I + \Sigma_1^2 + F^\top(I+\Sigma_2^2)F \preceq I + \Sigma_1^2 + \bar{\lambda} I \quad \Leftrightarrow \quad (I + \Sigma_1^2 + F^\top(I+\Sigma_2^2)F)^{-1} \succeq (I + \Sigma_1^2 + \bar{\lambda} I)^{-1},$$

so the relation is implied by semidefinite scaling. Put $\Lambda = (I + \Sigma_1^2 + \bar{\lambda} I)$ and note that the second diagonal block of (11) asserts $\Sigma_2 F \Lambda^{-1} F^\top \Sigma_2 \preceq I$, then

$$\begin{aligned}
S := (I+\Sigma^2)^{\frac{1}{2}}(I+\Sigma P_\Omega\Sigma)^{-1}(I+\Sigma^2)^{\frac{1}{2}} &\overset{(10)(11)}{\preceq} \\
&\preceq (I+\Sigma^2)^{\frac{1}{2}} \begin{bmatrix} I - \Sigma_1\Lambda^{-1}\Sigma_1 & -\Sigma_1\Lambda^{-1}F^\top\Sigma_2 \\ -\Sigma_2 F\Lambda^{-1}\Sigma_1 & I - \Sigma_2 F\Lambda^{-1}F^\top\Sigma_2 \end{bmatrix} (I+\Sigma^2)^{\frac{1}{2}} \\
&\preceq (I+\Sigma^2)^{\frac{1}{2}} \begin{bmatrix} (\Lambda - \Sigma_1^2)\Lambda^{-1} & -\Sigma_1\Lambda^{-1}F^\top\Sigma_2 \\ -\Sigma_2 F\Lambda^{-1}\Sigma_1 & I \end{bmatrix} (I+\Sigma^2)^{\frac{1}{2}} \quad := \bar{S}.
\end{aligned}$$

Employing [13, Prop. 8.3] now results in

$$\lambda_{\max}(S) \leq \lambda_{\max}(\bar{S}) \leq \|(I+\Sigma_1^2)(\Lambda - \Sigma_1^2)\Lambda^{-1}\| + \|I + \Sigma_2^2\|.$$

The last term evaluates to $\lambda_{\max}(I+\Sigma_2^2) = 1 + \sigma_{k+1}^2$. For the second last term substituting in the

definitions of Λ and $\bar{\lambda}$ yields

$$\begin{aligned}
&\|(I+\Sigma_1^2)(\Lambda-\Sigma_1^2)\Lambda^{-1}\| = \\
&= \quad (1+\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2)\cdot \max_{i\in\{1,\ldots,k\}} \frac{1+\sigma_i^2}{1+\sigma_i^2+\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2} \\
&= \quad \max_{i\in\{1,\ldots,k\}} \frac{1+\sigma_i^2+\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2+\sigma_i^2\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2}{1+\sigma_i^2+\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2} \\
&\leq \quad 1+\min\{\sigma_1^2, \|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2\}.
\end{aligned}$$

■

The current bound falls somewhat short of expectation because of the identity in $\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|^2$. By Theorem 1 and $I_n \preceq I_n + \Sigma P_\Omega \Sigma \preceq I_n + \Sigma^2$, the use of projectors will never result in condition numbers larger than $1+\sigma_1^2$, so the influence of the dimension seems to be too dominant in this. Maybe a better bound is achievable by a more sophisticated argument.

The deterministic bound allows to also make use of the probabilistic bounds on $\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\|$ given in [13]. These shed some light on the advantage of employing oversampling by p additional random vectors in Ω. In our application, oversampling corresponds to computing the singular values of $D^{-\frac{1}{2}}V\mathbf{P}_\Omega$ for $k+p$ columns in order to get better control on the k largest singular values of $D^{-\frac{1}{2}}V$ by the preconditioner $H_{\mathbf{P}_\Omega}$.

Theorem 7 *In the setting of Theorem 6 let Ω be drawn as a standard Gaussian matrix. Then*

$$\begin{aligned}
&\mathbb{E}\,\lambda_{\max}\big((I+\Sigma^2)^{\frac{1}{2}}(I+\Sigma P_\Omega\Sigma)^{-1}(I+\Sigma^2)^{\frac{1}{2}}\big) \leq \\
&\leq \quad 2+\sigma_{k+1}^2+\left(\sqrt{\frac{k}{p-1}}(1+\sigma_{k+1}^2)+\frac{e\sqrt{k+p}}{p}\big(\sum_{i=k+1}^{n}(1+\sigma_i^2)\big)^{\frac{1}{2}}\right)^2.
\end{aligned}$$

Furthermore, if $p\geq 4$ then for all $u,t\geq 1$ the probability for

$$\begin{aligned}
&\lambda_{\max}\big((I+\Sigma^2)^{\frac{1}{2}}(I+\Sigma P_\Omega\Sigma)^{-1}(I+\Sigma^2)^{\frac{1}{2}}\big) > \\
&> \quad 2+\sigma_{k+1}^2+\left(t\big[\sqrt{\tfrac{3k}{p+1}}+u\tfrac{e\sqrt{k+p}}{p+1}\big](1+\sigma_{k+1}^2)+t\tfrac{e\sqrt{k+p}}{p+1}\big(\sum_{i=k+1}^{n}(1+\sigma_i^2)\big)^{\frac{1}{2}}\right)^2
\end{aligned}$$

is at most $2t^{-p}+e^{-u^2/2}$.
The same bounds hold for the condition number κ_{P_Ω}.

Proof. A central and complex step in [13, proof of Th. 10.2] is to establish the relation

$$\mathbb{E}\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\| \leq \sqrt{\frac{k}{p-1}}\|(1+\Sigma_2^2)^{\frac{1}{2}}\| + \frac{e\sqrt{k+p}}{p}\|(I+\Sigma_2^2)^{\frac{1}{2}}\|_F$$

which directly yields the bound on the expected value via Theorem 6.

Likewise, in [13, proof of Th. 10.8] the authors derive for $p\geq 4$ and $u,t\geq 1$

$$\begin{aligned}
\mathbb{P}\Big\{\|(I+\Sigma_2^2)^{\frac{1}{2}}\Omega_2\Omega_1^\dagger\| > \|(I+\Sigma_2^2)^{\frac{1}{2}}\|\sqrt{\tfrac{3k}{p+1}}\cdot t + \|(I+\Sigma_2^2)^{\frac{1}{2}}\|_F\tfrac{e\sqrt{k+p}}{p+1}\cdot t \\
+\|(I+\Sigma_2^2)^{\frac{1}{2}}\|\tfrac{e\sqrt{k+p}}{p+1}\cdot ut\Big\} \leq 2t^{-p}+e^{-u^2/2}.
\end{aligned}$$

■

Again, the presence of the identity in the deterministic bound of Theorem 6 has a major impact also in these probabilistic bounds. Indeed, one would hope that a better deterministic bound helps to prove stronger decay.

Without some a priori knowledge of the singular values of $D^{-\frac{1}{2}}V \in \mathbb{R}^{m\times n}$ it is difficult to determine a suitable number of columns for Ω, i. e., a suitable dimension of the random subspace. For huge m the Johnson-Lindenstrauss result as presented in [8] suggests that for k at most $4(\varepsilon^2/2 - \varepsilon^3/3)^{-1} \ln m$ a suitably chosen random $\Omega \in \mathbb{R}^{n\times k}$ results in a distortion of $1 \pm \varepsilon$ with sufficiently high probability. This roughly translates to that each matrix element of $D^{-\frac{1}{2}}VV^\top D^{-\frac{1}{2}}$ and $D^{-\frac{1}{2}}V\Omega\Omega^\top V^\top D^{-\frac{1}{2}}$ differs by at most this factor. When considering the sizes of m aimed for here — the dimension of the design space will be a few thousands to a few hundred thousands — this number is still too large for efficient computations even for a moderate $\varepsilon = 0.1$. Indeed, the burden of forming the preconditioner and of applying it would exceed the gain by far. [13] propose an algorithmic variant for identifying a significant drop in singular values, but this requires successive matrix vector multiplications and these are quite costly in practice. In preliminary experiments with a number of tentative randomized variants, those relating the number of columns to the number of matrix-vector multiplications of the previous solve seemed reasonable. It will turn out, however, that even the cost of this is too high and the gain too small in comparison to the deterministic approach of the next section. The latter appears to capture the important directions quite well and offers better possibilities to exploit structural properties of the data. Due to the rather clear superiority of the deterministic approach, the numerical experiments of Section 4 will only present results for one particular randomized variant that performed best among the tentative versions. It attempts to identify the most relevant subspace by storing and extending the important directions generated in the previous round. For completeness and reproducability, its details are given in Alg. 8, but in view of the rather disencouraging results we refrain from further discussions.

Algorithm 8 (Randomized subspace selection forming $\hat{V} = VP_\Omega$)
Input: *$V \in \mathbb{R}^{m\times n}$, $D \in \mathbb{S}^n_{++}$, previous relevant subspace $P_{old} \in \mathbb{R}^{n\times \underline{k}}$ (initially $\underline{k} = 0$), previous number of multiplications n_{mult}, previous $\hat{k}$ of Alg. 3*
Output: *$\hat{V}$ (and stores P_{old})*
1. *If ($\underline{k} = 0$) then*
   (a) *set $k = \min\{n, 3 + 2\hat{k}, \lceil\sqrt{n_{mult}\frac{n_{mult}+n}{4}} - \frac{n_{mult}}{2}\rceil\}$,*
   (b) *generate a standard Gaussian $\Omega \in \mathbb{R}^{n\times k}$ and set $P_\Omega \leftarrow \Omega$,*
   else
   (a) *set $k_+ = \max\{3, \lfloor\frac{\sqrt{n_{mult}}}{2}\rfloor - \underline{k}\}$,*
   (b) *generate a standard Gaussian $\Omega \in \mathbb{R}^{n\times k_+}$ and set $P_\Omega \leftarrow [P_{old}, \Omega]$.*
2. *Orthonormalize P_Ω, reset k to its number of columns, set $\hat{V} = VP_\Omega$.*
3. *Compute eigenvalue decomposition $\hat{V}^\top D^{-1}\hat{V} = P\,\mathrm{Diag}(\hat{\lambda}_1 \geq \cdots \geq \hat{\lambda}_k)P^\top$.*
4. *Compute threshold $\bar{\lambda} = \max\{10, e^{\frac{1}{10}\ln\hat{\lambda}_1 - \frac{9}{10}\ln\hat{\lambda}_k}\}$ (enforce $\hat{\lambda}_k > 0$).*
5. *Set $\underline{k} \leftarrow \min\{k, \max\{3, i > 3 \colon \hat{\lambda}_i > \bar{\lambda}\}\}$ and set $P_{old} \leftarrow P_\Omega P_{\bullet,[1,\ldots,\underline{k}]}$.*
6. *Return $\hat{V}$.*

3.2 A Deterministic Subspace Selection Approach

In the conic bundle method, $H = D + VV^\top$ of Theorem 2 is of the form described in (8). An inspection of the column blocks of this V suggests to concretize the bound of Theorem 2 for interior point related applications as follows.

Theorem 9 *Given $D \in \mathbb{S}^m_{++}$ and $B \in \mathbb{R}^{n\times m}$, let $X \in \mathbb{S}^n_+$ have eigenvalue decomposition $X = [P_\ell, P_s]\left[\begin{smallmatrix}\Lambda_\ell & 0\\ 0 & \Lambda_s\end{smallmatrix}\right][P_\ell, P_s]^\top$ with $[P_\ell, P_s]^\top[P_\ell, P_s] = I_n$ and diagonal $\Lambda_\ell \in \mathbb{S}^k_+$, $\Lambda_s \in \mathbb{S}^{n-k}_+$. Put $V = B^\top X^{\frac{1}{2}}$. For $H = D + VV^\top$ and preconditioner $H_{P_\ell} = D + VP_\ell P_\ell^\top V^\top$ the condition number is bounded by*

$$\kappa_{P_\ell} \leq 1 + \sum_{i=1}^{n-k}(\Lambda_s)_{ii}\|B^\top(P_s)_{\bullet,i}\|^2_{D^{-1}} \leq 1 + (n-k)\bar{\rho}\bar{\beta}^2.$$

where $\bar{\rho} = \max_{i=1,\ldots,n-k}(\Lambda_s)_{ii}$ and $\bar{\beta} = \max_{i=1,\ldots,n}\|B_{i,\bullet}\|_{D^{-1}}$.

Proof. We show $\lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} V P_s P_s^\top V^\top H_{P_\ell}^{-\frac{1}{2}}) \leq \sum_{i=1}^{n-k} (\Lambda_s)_{ii} \|B^\top (P_s)_{\bullet,i}\|_{D^{-1}}^2$, then the statement follows by Theorem 2. Note that $V P_s P_s^\top V^\top = B^\top P_s \Lambda_s P_s^\top B$. Furthermore, if $\lambda \geq 0$ satisfies $B^\top B \preceq \lambda D$, it also satisfies $B^\top B \preceq \lambda(D + V P_\ell P_\ell^\top V^\top)$, therefore

$$\begin{aligned}
\lambda_{\max}(H_{P_\ell}^{-\frac{1}{2}} V P_s P_s^\top V^\top H_{P_\ell}^{-\frac{1}{2}}) &\leq \lambda_{\max}(D^{-\frac{1}{2}} B^\top P_s \Lambda_s P_s^\top B D^{-\frac{1}{2}}) \\
&\leq \operatorname{tr}(D^{-\frac{1}{2}} B^\top P_s \Lambda_s P_s^\top B D^{-\frac{1}{2}}) \\
&= \operatorname{tr}(\Lambda_s P_s^\top B D^{-1} B^\top P_s) \\
&= \sum_{i=1}^{n-k} (\Lambda_s)_{ii} \|B^\top (P_s)_{\bullet,i}\|_{D^{-1}}^2 \leq 1 + (n-k)\bar{\rho}\bar{\beta}^2.
\end{aligned}$$

■

Note, the proof weakens H_{P_ℓ} to D, so the bound cannot be expected to be strong. Yet, it provides a good rule of thumb on which columns of $D^{-\frac{1}{2}} B X^{\frac{1}{2}}$ should be included, namely those with large value $\Lambda_{ii} \|B^\top (P_s)_{\bullet,i}\|_{D^{-1}}^2$.

In interior point methods the spectral decomposition and the size of the eigenvalues of X of Theorem 9 strongly depend on the current iteration, in particular on the value of the barrier parameter μ. Therefore it is worth to set up a new preconditioner for each new KKT system. In order to do so in a computationally efficient way, the following dynamic selection heuristic for P_ℓ with respect to V of (8) tries to either pick columns of V directly by including unit vectors in P_ℓ or to at least take linear combinations of few columns of V in order to reduce the cost of matrix-vector multiplications and to preserve potential structural proprieties. So instead of forming P_ℓ, the heuristic builds $\hat{V} = V P_\ell$ directly by appending (linear combinations of selected) columns of V to $\hat{V}$. Also, it will often only employ approximate eigenvalues λ_i and eigenvectors p_i of the X described in Theorem 9. Generally, it will include those in $\hat{V}$ for which an estimate of $\lambda_i \|B^\top p_i\|_{D^{-1}}^2$ exceeds a given bound $\underline{\rho}$. In order to reduce the number of matrix vector multiplications, $\|B^\top p_i\|_{D^{-1}}^2$ will only be computed for those p_i with $\big(\sum_{j=1}^n (p_i)_j^2 \|(B^\top)_{\bullet,j}\|_{D^{-1}}\big)^2 \geq \underline{\rho}$ where the column norms of $B^\top$ are precomputed for each KKT systems. The implementation uses $\underline{\rho} = 10$. Next the selections are explained by going through V of (8) step by step for each of its three column groups V_H, $A^\top D_w^{\frac{1}{2}}$ and $B^\top \mathfrak{X}_t^{\frac{1}{2}} (I - \frac{\mathfrak{X}_t^{\frac{1}{2}} \mathbb{1}_t (\mathfrak{X}_t^{\frac{1}{2}} \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t})^{\frac{1}{2}}$. Concerning the third group, it will become clear in the discussion of the semidefinite part that in practice it is advantageous to replace the square root $\mathfrak{X}_t^{\frac{1}{2}}$ in the factorization of $\mathfrak{X}_t$ by a more general, possibly nonsymmetric factorization $\mathfrak{X}_t = \mathfrak{F}_t \mathfrak{F}_t^\top$. The matrix $\mathfrak{F}_t$ will have the same block structure and leads to a similar rank one correction by the transformed trace vector $\mathfrak{F}_t^\top \mathbb{1}_t$,

$$\mathfrak{X}_t - \frac{\mathfrak{X}_t \mathbb{1}_t (\mathfrak{X}_t \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t} = \mathfrak{F}_t (I - \frac{\mathfrak{F}_t^\top \mathbb{1}_t (\mathfrak{F}_t^\top \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{F}_t \mathfrak{F}_t^\top \mathbb{1}_t})^{\frac{1}{2}} (I - \frac{\mathfrak{F}_t^\top \mathbb{1}_t (\mathfrak{F}_t^\top \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{F}_t \mathfrak{F}_t^\top \mathbb{1}_t})^{\frac{1}{2}} \mathfrak{F}_t^\top.$$

Algorithm 10 (Deterministic column selection heuristic forming $\hat{V}$)
Input: *D_H, V_H, D_y, A, D_w, B, $\mathfrak{X}_t$, ζ, σ specifying D and $V \in \mathbb{R}^{m \times n}$ of (8)*
Output: *$\hat{V} \in \mathbb{R}^{m \times n'}$ for some $n' \leq n$ with $\hat{V} = V P_\ell$, $P_\ell^\top P_\ell = I_{n'}$.*

1. *Initialize $\hat{V} \leftarrow 0 \in \mathbb{R}^{m \times 0}$, $\underline{\rho} := \frac{10}{m} \operatorname{tr} D_H$*
2. *Find $\mathcal{J}_{V_h} = \{j \colon \|(V_H)_{\bullet,j}\|_{D^{-1}}^2 \geq \underline{\rho}\}$ and set $\hat{V} \leftarrow [\hat{V}, (V_H)_{\bullet,\mathcal{J}_{V_H}}]$.*
3. *Find $\mathcal{J}_A = \{j \colon (D_w)_{jj} \|(A^\top)_{\bullet,j}\|_{D^{-1}}^2 \geq \underline{\rho}\}$ and set $\hat{V} \leftarrow [\hat{V}, (A^\top D_w)_{\bullet,\mathcal{J}_A}]$.*
4. *Compute $\mathfrak{F}_t^\top \mathbb{1}_t$, $\eta = \zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{F}_t \mathfrak{F}_t^\top \mathbb{1}_t$, $B^\top \mathfrak{F}_t^\top \mathbb{1}_t$ and for each conic diagonal block of $\mathfrak{X}_t$ call append_"cone"_columns($\hat{V}$) with corresponding parameters.*
5. *Return $\hat{V}$.*

The first group of columns $V_H \in \mathbb{R}^{m \times h_H}$ matches, in the notation of Theorem 9, (a subblock of) $B^\top = V_H$ and (a diagonal block) $X = I_{h_H}$. The heuristic appends those columns j to $\hat{V}$ that satisfy $\|(V_H)_{\bullet,j}\|_{D^{-1}}^2 \geq \underline{\rho}$.

For the second group of columns $A^\top D_w^{\frac{1}{2}}$, Theorem 9 applies to $B^\top = A^\top$ and $X = D_w = \mathrm{Diag}(d_1, \ldots, d_{h_A})$. Thus, column j is appended to $\hat{V}$ if $d_j \|(A^\top)_{\bullet,j}\|^2_{D^{-1}} \geq \underline{\rho}$.

With the comment above regarding $\mathfrak{F}_t$, the third column group is formed by a term $B^\top \mathfrak{F}_t (I - \frac{\mathfrak{F}_t^\top \mathbb{1}_t (\mathfrak{F}_t^\top \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t})^{\frac{1}{2}}$ for each cutting model (we assume just one here). With respect to Theorem 9, B is just right and X is the positive (semi-)definite matrix $\mathfrak{X}_t - \frac{\mathfrak{X}_t \mathbb{1}_t (\mathfrak{X}_t \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t} = \mathfrak{F}_t (I - \frac{\mathfrak{F}_t \mathbb{1}_t (\mathfrak{F}_t \mathbb{1}_t)^\top}{\zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{F}_t \mathfrak{F}_t^\top \mathbb{1}_t}) \mathfrak{F}^\top$. Recall that $\mathfrak{X}_t$ is a block diagonal matrix with the structure of the diagonal blocks governed by the linearization of the perturbed complementarity conditions of the various cones. The overarching rank one modification by $\mathfrak{F}_t \mathbb{1}_t$ couples the blocks within the same cutting model and reappears in some form in each block together with $\eta = \zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t$. Observe that with $\|\mathfrak{F}_t^\top \mathbb{1}_t\|^2 = \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t$

$$(I - \tfrac{\mathfrak{F}_t^\top \mathbb{1}_t (\mathfrak{F}_t^\top \mathbb{1}_t)^\top}{\eta})^{\frac{1}{2}} = I - \big(1 - \tfrac{\sqrt{\zeta^{-1}\sigma}}{\sqrt{\eta}}\big) \tfrac{\mathfrak{F}_t^\top \mathbb{1}_t}{\|\mathfrak{F}_t^\top \mathbb{1}_t\|} \tfrac{(\mathfrak{F}_t^\top \mathbb{1}_t)^\top}{\|\mathfrak{F}_t^\top \mathbb{1}_t\|}.$$

For each column p of P_ℓ computing $B^\top \mathfrak{F}_t (I - \frac{\mathfrak{F}_t^\top \mathbb{1}_t (\mathfrak{F}_t^\top \mathbb{1}_t)^\top}{\eta})^{\frac{1}{2}} p$ splits into

$$B^\top \mathfrak{F}_t p - \big\langle \mathfrak{F}_t^\top \mathbb{1}_t, p \big\rangle \tfrac{1}{\mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t} \big(1 - \tfrac{\sqrt{\zeta^{-1}\sigma}}{\sqrt{\eta}}\big) B^\top \mathfrak{X}_t \mathbb{1}_t.$$

Thus, by keeping the support of p restricted to single blocks, the proper column computations can be kept restricted to the respective block. This also holds for the coefficient $\big\langle \mathfrak{F}_t^\top \mathbb{1}_t, p \big\rangle$. The overarching vector $B^\top \mathfrak{X}_t \mathbb{1}_t$ needs to be evaluated only once and can be added to the columns afterwards. The latter step only requires the respective coefficients but not the vectors of P_ℓ. This allows to speed up the process of forming $\hat{V}$ considerably. Therefore, when forming the conceptional P_ℓ in the heuristic, the influence of $\mathfrak{F}_t^\top \mathbb{1}_t$ on eigenvalues and eigenvectors of the blocks will mostly be considered as restricted to each single block. Next the actual selection procedure is described for $\mathfrak{X}_t$ blocks corresponding to cones $\mathbb{R}^h_+$ (Alg. 11) and $\mathbb{S}^h_+$ (Alg. 13) with Nesterov-Todd-scaling [26, 30].

Algorithm 11 (append_$\mathbb{R}^h_+$_columns($\hat{V}$))
Input: *column indices $J \in \mathbb{N}^h$ and $x \circ z^{-1}$ of this block in $\mathfrak{X}_t$, $B^\top_{\bullet,J}$, $\mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t$, $B^\top \mathfrak{X}_t \mathbb{1}_t$, $\eta = \zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t$, D, threshold $\underline{\rho}$.*
Output: *updated $\hat{V}$.*
1. For each $i = 1, \ldots, h$ with $(\frac{x_i}{z_i} - \frac{1}{\eta}\frac{x_i^2}{z_i^2}) \|(B^\top)_{\bullet,J(i)}\|^2_{D^{-1}} \geq \underline{\rho}$ set

$$\alpha \leftarrow \sqrt{\tfrac{x_i}{z_i}} \tfrac{1}{\mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t} \big(1 - \tfrac{\sqrt{\zeta^{-1}\sigma}}{\sqrt{\eta}}\big),$$

$$\hat{b}_i = \sqrt{\tfrac{x_i}{z_i}} (B^\top)_{\bullet,J(i)} - \alpha B^\top \mathfrak{X}_t \mathbb{1}_t,$$

and if $\|\hat{b}_i\|^2_{D^{-1}} > \underline{\rho}$ set $\hat{V} \leftarrow [\hat{V}, \hat{b}_i]$.

For Alg. 11 consider, within the cone specified by t, a block with indices $J \in \mathbb{N}^h$ representing a nonnegative cone $\mathbb{R}^h_+$ with primal dual pair (x, z). The corresponding "diagonal block" in $\mathfrak{X}_t$ is of the form $\mathrm{Diag}(x \circ z^{-1})$ and for $\mathfrak{F}_t$ it is $\mathrm{Diag}(x \circ z^{-1})^{\frac{1}{2}}$. The relevant part of the trace vector $\mathfrak{X}_t \mathbb{1}_t$ reads $\mathrm{Diag}(x \circ z^{-1})\mathbb{1} = x \circ z^{-1}$. Considering the influence of the trace vector as restricted to this block alone gives $\mathrm{Diag}(x \circ z^{-1}) - \frac{1}{\eta}(x \circ z^{-1})(x \circ z^{-1})^\top$ with the correct overarching $\eta = \zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t$. The eigenvectors to large eigenvalues of this matrix have their most important coordinates associated with the largest diagonal entries. The heuristic appends the columns $B^\top \mathfrak{X}_t^{\frac{1}{2}} (I - \frac{\mathfrak{X}_t^{\frac{1}{2}} \mathbb{1}_t (\mathfrak{X}_t^{\frac{1}{2}} \mathbb{1}_t)^\top}{\eta})^{\frac{1}{2}} e_{J(i)}$ to $\hat{V}$ for those $e_{J(i)}$ with $(\frac{x_i}{z_i} - \frac{1}{\eta}\frac{x_i^2}{z_i^2}) \|(B^\top)_{\bullet,J(i)}\|^2_{D^{-1}} \geq \underline{\rho}$.

Note that in interior point methods $x_i z_i \approx \mu$ for barrier parameter $\mu \searrow 0$ and $x_i \to x_i^{opt}$, $z_i \to z_i^{opt}$. Due to $\eta \geq \frac{x_i}{z_i}$ with η mostly much larger, the estimated value roughly behaves like $\frac{x_i^2}{\mu} \|(B^\top)_{\bullet,J(i)}\|^2_{D^{-1}}$ and, indeed, by experience it seems that columns are almost exclusively

included only for active $x_i^{opt} > 0$ and only as μ gets small enough. When computing high precision solutions with small μ, the rank of the preconditioner can thus be expected to match the number of active subgradients in the cutting model. Theorem 9 suggests that in iterative methods these columns have to be included in some form in order to obtain reliable convergence behavior.

Alg. 13 below deals with a positive semidefinite cone $\mathbb{S}_+^h$ with Nesterov-Todd-scaling. For the current purposes it suffices to know that the diagonal block of $\mathfrak{X}_t$ indexed by appropriate $J \in \mathbb{N}^{\binom{h+1}{2}}$ is of the form $W \otimes_s W$ for a positive definite $W \in \mathbb{S}_{++}^h$; see [30] for its efficient computation and for an appendix of convenient rules for computing with symmetric Kronecker products. The next result derives the eigenvectors and eigenvalues when considering the rank one correction restricted to this block.

Lemma 12 *Let $W = P_W \Lambda_W P_W^\top$ with $\Lambda_W = \operatorname{Diag}(\lambda_1^W \geq \cdots \geq \lambda_h^W > 0)$ and $P_W^\top P_W = I_h$, $P_W = [w_1, \ldots, w_h]$. Furthermore let $U = \Lambda_W^2 - \frac{1}{\eta}(\Lambda_W^2 \mathbb{1})(\Lambda_W^2 \mathbb{1})^\top$ have eigenvalue decomposition $U = P_U \Lambda_U P_U^\top$ with $P_U^\top P_U = I_h$. The eigenvalues of $W \otimes_s W - \frac{1}{\eta}\big((W \otimes_s W)\operatorname{svec}(I_n)\big)\big((W \otimes_s W)\operatorname{svec}(I_h)\big)^\top$ are $\lambda_i^U = (\Lambda_U)_{ii}$ with eigenvectors $\sum_{j=1}^h (P_U)_{ji} \operatorname{svec}(w_j w_j^\top)$ for $i = 1, \ldots, h$ and $\lambda_i^W \lambda_j^W$ with eigenvectors $\frac{1}{\sqrt{2}} \operatorname{svec}(w_i w_j^\top + w_j w_j^\top)$ for $1 \leq i < j \leq h$.*

Proof. By [2] the eigenvalues of $(W \otimes_s W)$ are $\lambda_i^W \lambda_j^W$ for $1 \leq i \leq j \leq h$ with orthonormal eigenvectors

$$\begin{array}{ll} w_{ii} := \operatorname{svec}(w_i w_i^\top) & \text{for } 1 \leq i = j \leq h, \\ w_{ij} := \frac{1}{\sqrt{2}} \operatorname{svec}(w_i w_j^\top + w_j w_i^\top) & \text{for } 1 \leq i < j \leq h. \end{array} \tag{12}$$

To see this *e.g.* for $i < j$ observe $w_{ij}^\top w_{ij} = \frac{1}{2}[2 \langle w_i w_j^\top, w_i w_j^\top \rangle + 2 \langle w_i w_j^\top, w_j w_i^\top \rangle]$ and $(W \otimes_s W) w_{ij} = \frac{1}{\sqrt{2}} \operatorname{svec}(W w_i w_j^\top W + W w_j w_i^\top W) = \lambda_i^W \lambda_j^W w_{ij}$.

From $(W \otimes_s W)\operatorname{svec}(I_h) = \operatorname{svec}(W^2) = \sum_{i=1}^h (\lambda_i^W)^2 w_{ii}$ one obtains

$$\begin{array}{ll} \operatorname{svec}(W^2)^\top w_{ii} = (\lambda_i^W)^2 & \text{for } i = 1, \ldots, h, \\ \operatorname{svec}(W^2)^\top w_{ij} = 0 & \text{for } 1 \leq i < j \leq h. \end{array}$$

The eigenvector-sorting $P_{\otimes_s} = [w_{11}, w_{22}, \ldots, w_{hh}, w_{12}, w_{13}, \ldots, w_{h-1,h}]$ gives

$$W \otimes_s W - \frac{1}{\eta} \operatorname{svec}(W^2) \operatorname{svec}(W^2)^\top = P_{\otimes_s} \begin{bmatrix} U & 0 & \cdots & 0 \\ 0 & \lambda_1^W \lambda_2^W & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_{h-1}^W \lambda_h^W \end{bmatrix} P_{\otimes_s}^\top$$

The result now follows by direct computation. ∎

For semidefinite blocks, numerical experience indicates that it is indeed worth to determine the eigenvalue decomposition of U as in Lemma 12. Finding the eigenvalues and eigenvectors roughly requires the same amount of work as forming W and is of no concern. With $J \in \mathbb{N}^{\binom{h+1}{2}}$ denoting the column indices of this block within $B^\top$, columns to corresponding eigenvectors are computed by $(\Lambda_U^{\frac{1}{2}})_{ii} \cdot (B^\top)_{\bullet,J} \sum_{j=1}^h (P_U)_{ji} w_{jj}$ or $\sqrt{\lambda_i^W \lambda_j^W}(B^\top)_{\bullet,J} w_{ij}$. This involves linear combinations of $\binom{h+1}{2}$ columns and is computationally expensive if the order h of W gets large. Indeed, when testing all columns by their correct norms $\|(B^\top)_{\bullet,J} w_{ij}\|_{D^{-1}}^2$, too much time is spent in forming the preconditioner. Therefore the heuristic Alg. 13 first selects candidate eigenvectors to use for P_ℓ via the rough estimate $\sum_{i=1}^{\binom{h}{2}} (w_{ij})_i^2 \|(B^\top)_{\bullet,J(i)}\|_{D^{-1}}^2 = \|w_{ij}\|_{\operatorname{Diag}(BD^{-1}B^\top)_J}^2$. For the selected eigenvectors it then computes the precise values after the following transformation that is only seemingly involved.

In order to also account for the possibly overarching contribution of $\mathfrak{F}_t \mathbb{1}_t$ it is advantageous to find a representation equivalent to $B^\top X^{\frac{1}{2}} P_\ell$ with orthonormal columns in P_ℓ as in Theorem 9 for a suitable factorization of X other than its square root. For this, let $V_W = P_W \Lambda_W^{\frac{1}{2}}$, then $W \otimes_s W = (V_W \otimes_s V_W)(V_W^\top \otimes_s V_W^\top)$. Because $(V_W^\top \otimes_s V_W^\top) \operatorname{svec} I = \operatorname{svec}(V_W^\top V_W) = \operatorname{svec} \Lambda_W$ and

$(V_W \otimes_s V_W) = (P_W \otimes_s P_W)(\Lambda_W^{\frac12} \otimes_s \Lambda_W^{\frac12})$, the notation of Lemma 12 and its proof allows to rephrase the semidefinite block of $\mathfrak{X}_t - \frac{\mathfrak{X}_t \mathbb{1}_t (\mathfrak{X}_t \mathbb{1}_t)^\top}{\eta}$ as

$$\begin{aligned}
&W \otimes_s W - \frac{(W \otimes_s W)\operatorname{svec}(I)\operatorname{svec}(I)^\top (W \otimes_s W)}{\eta} = \\
&= (V_W \otimes_s V_W)(I - \frac{\operatorname{svec}(\Lambda_W)\operatorname{svec}(\Lambda_W)^\top}{\eta})(V_W \otimes_s V_W)^\top \\
&= (P_W \otimes_s P_W)(\Lambda_W^{\frac12} \otimes_s \Lambda_W^{\frac12})(I - \frac{\operatorname{svec}(\Lambda_W)\operatorname{svec}(\Lambda_W)^\top}{\eta})(\Lambda_W^{\frac12} \otimes_s \Lambda_W^{\frac12})(P_W \otimes_s P_W)^\top \\
&= P_{\otimes_s} \begin{bmatrix} U & 0 & \cdots & 0 \\ 0 & \lambda_1^W \lambda_2^W & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \lambda_{h-1}^W \lambda_h^W \end{bmatrix} P_{\otimes_s}^\top \\
&= (P_W \otimes_s P_W) F F^\top (P_W \otimes_s P_W)^\top,
\end{aligned}$$

where

$$F = (\Lambda_W^{\frac12} \otimes_s \Lambda_W^{\frac12})(I - \frac{\operatorname{svec}(\Lambda_W)\operatorname{svec}(\Lambda_W)^\top}{\eta})^{\frac12}.$$

This suggests to put $V = B^\top P_{\otimes_s} F$ and to derive the columns corresponding to P_ℓ via the singular value decomposition of $F = Q_F \Sigma_F P_F^\top$. Lemma 12 provides the squared singular values Σ_F^2 and the eigenvectors give the left-singular vectors in Q_F. For $e_{ij} := \frac{1}{\sqrt2}\operatorname{svec}(e_i e_j^\top + e_j e_i^\top)$ there holds $\operatorname{svec}(\lambda_W)^\top e_{ij} = 0$, so the right-singular vectors corresponding to $\sqrt{\lambda_i \lambda_j}$ read $(P_F)_{\bullet,ij} = e_{ij}$. The remaining right-singular vectors of P_F may be computed via $P_F = F^\top Q_F \Sigma_F^{-1}$. In this it is sufficient and convenient to consider only the U block, i. e., the support restricted to the ii-coordinates. Denote the columns of $P_U = [u_1, \ldots, u_h]$ in Lemma 12 by u_j for $j = 1, \ldots, h$, then the corresponding right-singular vectors $u_j^F \in \mathbb{R}^h$ read for $\Lambda_W = \operatorname{Diag}(\lambda^W)$ and $\Lambda_U = \operatorname{Diag}(\lambda_1^U, \ldots, \lambda_h^U)$

$$\begin{aligned}
u_j^F &= (I - \tfrac{\lambda^W (\lambda^W)^\top}{\eta})^{\frac12} \Lambda_W \cdot u_j \cdot \tfrac{1}{\sqrt{\lambda_j^U}} \qquad (13)\\
&= \tfrac{1}{\sqrt{\lambda_j^U}} \left(\Lambda_W u_j - \frac{\sqrt{\eta} - \sqrt{\eta - \|\lambda^W\|^2}}{\sqrt{\eta}\|\lambda^W\|^2} \langle \lambda^W, \Lambda_W u_j \rangle \lambda^W \right).
\end{aligned}$$

By expanding the U block to the correct positions, the right-singular vector to singular value $\sqrt{\lambda_j^U}$ is $(P_F)_{\bullet,jj} = \operatorname{svec}(\operatorname{Diag}(u_j^F))$ for $j = 1, \ldots, h$.

With these preparations the selected semidefinite columns are appended to $\hat V$ as follows. First note that the semidefinite block with coordinates J of the factor $\mathfrak{F}_t$ is $(V_W \otimes_s V_W)$, which is non-symmetric in general. The transformed trace vector $\mathfrak{F}_t^\top \mathbb{1}_t$ reads $(\mathfrak{F}_t^\top \mathbb{1}_t)_J = (V_W^\top \otimes_s V_W^\top) \operatorname{svec} I = \operatorname{svec}(\Lambda_W)$. If column p_{ij}^F of P^F with $1 \le i \le j \le h$ is selected for P_ℓ by the heuristic, the column to be appended to $\hat V$ reads

$$(B^\top)_{\bullet,J}(V_W \otimes_s V_W) p_{ij}^F - \langle \operatorname{svec}\Lambda_W, p_{ij}^F \rangle \tfrac{1}{\mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t} \big(1 - \tfrac{\sqrt{\zeta^{-1}\sigma}}{\sqrt{\eta}}\big) B^\top \mathfrak{X}_t \mathbb{1}_t.$$

If the selected indices satisfy $i < j$, the vector p_{ij}^F is just $e_{ij} = \frac{1}{\sqrt2}\operatorname{svec}(e_i e_j^\top + e_j e_i^\top)$. By $(V_W \otimes_s V_W) e_{ij} = \frac{\sqrt{\lambda_i^W \lambda_j^W}}{\sqrt2}\operatorname{svec}(w_i w_j^\top + w_j w_i^\top) = \sqrt{\lambda_i^W \lambda_j^W} w_{ij}$ and $\langle \operatorname{svec}\Lambda_W, e_{ij} \rangle = 0$ the column computation simplifies to

$$\sqrt{\lambda_i^W \lambda_j^W}(B^\top)_{\bullet,J} w_{ij} = \sqrt{2\lambda_i^W \lambda_j^W}\big[w_i^\top \operatorname{svec}^{-1}([B^\top]_{k,J}) w_j\big]_{k=1,\ldots,n}.$$

Typically, several mixed eigenvectors w_{ij} have the same index i corresponding to a large value λ_i^W, so it quickly pays off to precompute $w_i^\top \operatorname{svec}^{-1}([B^\top]_{k,J})$ and to use these h-vectors for each

w_j. For ease of presentation this implementational detail is not described in Alg. 13. Also, this is not helpful for the non-mixed vectors $p^F_{jj} = \operatorname{svec}(\operatorname{Diag}(u^F_j))$, because

$$(V_W \otimes_s V_W) p^F_{jj} = (V_W \otimes_s V_W) \sum_{i=1}^{h} (u^F_j)_i \operatorname{svec} e_i e_i^\top = \sum_{i=1}^{h} (u^F_j)_i \lambda^W_i w_{ii}$$

consists of a linear combination over all w_{ii}. Fortunately only few of the non-mixed vectors are among those selected for preconditioning. The transformed trace vector coefficient for p^F_{jj} evaluates to $\langle \Lambda_W, \operatorname{Diag}(u^F_j) \rangle = \langle \lambda^W, u^F_j \rangle$. With this, the algorithm for appending semidefinite columns reads as follows.

Algorithm 13 (append_$\mathbb{S}^h_+$_columns($\hat{V}$))

Input: *column indices $J \in \mathbb{N}^{\binom{h+1}{2}}$ and Nesterov-Todd scaling matrix $W \succ 0$ of this block in $\mathfrak{X}_t$, $B^\top_{\bullet,J}$, $\mathbb{1}_t \mathfrak{X}_t \mathbb{1}_t$, $B^\top \mathfrak{X}_t \mathbb{1}_t$, $\eta = \zeta^{-1}\sigma + \mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t$, D, threshold $\underline{\rho}$*

Output: *updated $\hat{V}$.*

1. *Compute norms $\|(B^\top)_{\bullet,J(i)}\|_{D^{-1}}$, set $\hat{\rho} = \underline{\rho} / \max_{i=1,\dots,\binom{h+1}{2}} \|(B^\top)_{\bullet,J(i)}\|^2_{D^{-1}}$, compute eigenvalue decomposition $W = P_W \Lambda_W P_W^\top$, $\Lambda_W = \operatorname{Diag}(\lambda^W)$ with $\lambda^W_1 \geq \cdots \geq \lambda^W_h$, let w_{ij} be defined by (12).*
2. *If $(\lambda^W_1)^2 < \hat{\rho}$ do nothing and return $\hat{V}$.*
3. *Compute $U = \Lambda^2_W - \frac{1}{\eta}(\lambda^W)(\lambda^W)^\top$, eigenvalue decomposition $U = P_U \Lambda_U P_U^\top$, with $P = [u_1, \dots, u_h]$.*
4. *For each $\hat{\imath} = 1, \dots, h$ with $(\Lambda_U)_{\hat{\imath}\hat{\imath}} \geq \hat{\rho}$ do:*
   Compute $\hat{w}_{\hat{\imath}\hat{\imath}} = \sum_{i=1}^h (u_{\hat{\imath}})_i w_{ii}$ $(\in \mathbb{R}^{\binom{h+1}{2}})$.
   If $(\Lambda_U)_{\hat{\imath}\hat{\imath}} \sum_{j=1}^{\binom{h+1}{2}} (\hat{w}_{\hat{\imath}\hat{\imath}})^2_j \|(B^\top)_{\bullet,J(j)}\|^2_{D^{-1}} \geq \underline{\rho}$ then:
   (a) Compute $u^F_{\hat{\imath}}$ according to (13) and set

$$\alpha \leftarrow \langle \lambda^W, u^F_{\hat{\imath}} \rangle \tfrac{1}{\mathbb{1}_t^\top \mathfrak{X}_t \mathbb{1}_t} \big(1 - \tfrac{\sqrt{\zeta^{-1}\sigma}}{\sqrt{\eta}}\big),$$

$$\hat{b}_{\hat{\imath}\hat{\imath}} = (B^\top)_{\bullet,J} \sum_{i=0}^{h} (u^F_{\hat{\imath}})_i \lambda^W_i w_{ii} - \alpha B^\top \mathfrak{X}_t \mathbb{1}_t.$$

   (b) If $\|\hat{b}_{\hat{\imath}\hat{\imath}}\|^2_{D^{-1}} \geq \underline{\rho}$ set $\hat{V} \leftarrow [\hat{V}, \hat{b}_{\hat{\imath}\hat{\imath}}]$.
5. *For each $1 \leq \hat{\imath} < \hat{\jmath} \leq h$ with $\lambda^W_{\hat{\imath}} \lambda^W_{\hat{\jmath}} > \hat{\rho}$ do:*
   If $\sqrt{\lambda^W_{\hat{\imath}} \lambda^W_{\hat{\jmath}}} \sum_{j=1}^{\binom{h+1}{2}} (\hat{w}_{\hat{\imath}\hat{\jmath}})^2_j \|(B^\top)_{\bullet,J(j)}\|^2_{D^{-1}} \geq \underline{\rho}$ set

$$\hat{b}_{\hat{\imath}\hat{\jmath}} = \sqrt{\lambda^W_{\hat{\imath}} \lambda^W_{\hat{\jmath}}} (B^\top)_{\bullet,J} w_{\hat{\imath}\hat{\jmath}}$$

   and if $\|\hat{b}_{\hat{\imath}\hat{\jmath}}\|^2_{D^{-1}} \geq \underline{\rho}$ set $\hat{V} \leftarrow [\hat{V}, \hat{b}_{\hat{\imath}\hat{\jmath}}]$.
6. *Return $\hat{V}$.*

As for the linear case it can be argued that for small barrier parameter μ the number of selected columns corresponds at least to the order of the active submatrix in the cutting model. Thus if $\hat{h} \leq h$ eigenvalues of $X \in \mathbb{S}^h_+$ converge to positive values in the optimum, the heuristic will end up with selecting at least $\binom{\hat{h}+1}{2}$ columns once μ gets small.

For second order cones $\mathcal{Q}^h$ the structural properties of the arrow operator and the Nesterov-Todd-direction allow to restrict considerations to just two directions per cone for preconditioning, but as the computational experiments do not involve second order cones this will not be discussed here.

4 Numerical Experiments

The purpose of the numerical experiments is to explore and compare the behavior and performance of the pure and preconditioned iterative variants to the original direct solver on KKT instances that arise in the course of solving large scale instances by the conic bundle method.

It has to be emphasized that the experiments are by no means designed and intended to investigate the efficiency of the conic bundle method with internal iterative solver. Indeed, many aspects of the ConicBundle code [15] such as the cutting model selection routines, the path following predictor-corrector approach and the internal termination criteria have been tuned to work reasonably well with the direct solver. As the theory suggests and the results support, the performance of iterative methods depends more on the size of the active set than on the size of the model. Thus somewhat larger models might be better in connection with iterative solvers. Also, the predictor-corrector approach is particularly efficient if setting up the KKT system is expensive. For iterative methods with deterministic preconditioning this hinges on the cost of forming the preconditioner which gets expensive once the barrier parameter gets small. Furthermore iterative methods might actually profit from staying in a rather narrow neighborhood of the central path. Therefore many implementational decisions need to be reevaluated for iterative solvers. This is out of scope for this paper. Hence, the experiments only aim to highlight the relative performance of the solvers on sequences of KKT systems that currently arise in ConicBundle.

The experiments will report on the performance for three different instances: the first, denoted by MC, is a classical semidefinite relaxation of Max-Cut on a graph with 20000 nodes as described in [11, 19], the second, BIS, is a semidefinite Minimum-Bisection relaxation improved by dynamic separation of odd cycle cutting planes on the support of the Boeing instance KKT_traj33 giving a graph on 20006 nodes explained in [14], and the third, MMBIS, refers to a min-max-bisection problem problem shifting the edge weights so as to minimize a restricted maximum cut on a graph of 12600 nodes. All three have a single semidefinite cutting model which consists of a semidefinite cone with up to one nonnegative variable, so the model cone $\mathcal{S}_+^t$ of (2) typically has $t = (1, [], [h])$ for some $h \in \mathbb{N}$. In the Max-Cut instance the design variables are unconstrained, in the Bisection instance the design variables corresponding to the cutting planes are sign constrained (D_y is needed) and in the min-max-bisection problem some design variables have bounds and there are linear equality and inequality constraints (D_y, D_w and A appear). Throughout, the proximal term is a multiple of the identity for a dynamic weight, i. e., $H_k = u_k I$ with $u_k > 0$ controlled as in [16].

In each case ConicBundle is run with default settings for the internal constrained QP solver with direct KKT solver for the bundle subproblems. Whenever a new KKT-system arises, it is solved consecutively but independently on the same machine by

- (DS) the original direct solver,
- (IT) MINRES without preconditioning (the implementation follows [10]),
- (RP) MINRES with randomized preconditioning (Alg. 3 with Alg. 8),
- (DP) MINRES with deterministic preconditioning (Alg. 3 with Alg. 10).

Only the results of the direct solver are then used to continue the algorithm. Note, for nonsmooth optimization problems tiny deviations in the solution of the subproblem may lead to huge differences in the subsequent path of the algorithm. Therefore running the bundle method with different solvers would quickly lead to incomparable KKT systems.

The details of the direct solver DS are of little relevance at this point. Suffice it to say that its main work consists in Schur complementing the H and $\zeta^{-1}\sigma$ blocks of the KKT system (6) into the joined $\mathrm{Diag}(D_w^{-1}, \mathfrak{X}_t^{-1})$ block and factorizing this. In the Max-Cut setting (no D_y), the H block is constant throughout each bundle subproblem. In this case the Schur complement is precomputed once for each bundle subproblem — thus for several KKT systems — and this make this approach extremely efficient as long as the order h of the semidefinite model is small. Precomputation is no longer possible if D_y is needed which is the case in the two other instances. Finally, if A is also present, the system to be factorized in every iteration gets significantly larger. These differences motivated the choice of the instances and explain part of the strong differences in the performance of the solvers.

For Max-Cut and Bisection the iterative solver could exploit the positive definiteness of the system by employing conjugate gradients instead of MINRES. The min-max-bisection problem comprises equality constraints in A, so the system is no longer positive definite and conjugate gradients are not applicable. Employing MINRES for all three facilitates the comparison, in particular as MINRES seemed to perform numerically better on the other instances as well. MINRES computes the residual norm with respect to the inverse of the preconditioner and the implementation uses this norm for termination. To safeguard against effects due to the changes in this norm, the relative precision requirement is multiplied, in the notation of Alg. 3, by the factor $\big(\sqrt[m]{\prod_{i=1}^{\hat{k}}(1+\hat{\lambda}_i)^{-1}}\cdot \min_i(D^{-1})_i\big)^{\frac{1}{2}}$.

The results on the three instances will be presented in eight plots per instance. The first four compare all four solvers, the last four plots are devoted to information that is only relevant for iterative solvers, so DS will not appear in these.

1. Plot "time per subproblem (seconds)" gives for each of the four methods a box plot on the seconds (in logarithmic scale) required to solve the subproblems. For each subproblem this is the sum of the time required for initializing/forming and solving all KKT systems of this subproblem. This is needed, because in the case of Max-Cut instance MC, the direct solver DS forms the Schur complement of the H-block only once per subproblem and this is also accounted for here.
2. Plot "subproblem time (seconds) per iteration" displays the same cumulative time per subproblem in seconds (in logarithmic scale) for each successive iteration so that the development in solution time is aligned to the progress of the bundle method.
3. Plot "time per subproblem vs. bundle size" serves to highlight the dependence of the solution time on the size of the cutting model (number of rows of B). For this the subproblems are grouped in the bundle size ranges $(0, 50]$, $(50, 500]$, $(500, 1500]$, $(1500, \infty]$. Instead of these numbers, the actual maxima of the groups are shown in the bottom line of the plot.
4. Plot "time per subproblem vs. last μ" illustrates the dependence of the solution time on the last barrier parameter μ for which the subproblem has to be solved. Roughly this corresponds to the precision required for the subproblem. Results are presented for a subdivision of the subproblems into four groups of equal cardinality (up to integer division) sorted according to the μ value of their respective last KKT system. The minimum μ of each group is given in the bottom line of the plot.
5. Plot "time per KKT system (seconds)" compares exclusively the iterative methods on the KKT systems belonging to the four different ranges of the barrier parameter μ as collected over all subproblems. The first three box plots give the box plot statistics on the seconds (in logarithmic scale) spent in solving KKT systems for barrier parameter values $\mu \geq 100$, the next three for $100 > \mu \geq 1$, etc. Note, DS would require the same time for all KKT systems of the same subproblem, because its solution time does not depend on μ or the associated required relative precision which was set to $\min\{0.01 \cdot \mu, 10^{-6}\}$.
6. Plot "matrix vector multiplications per KKT system" shows box plots on the number of matrix-vector multiplications (in logarithmic scale) needed by MINRES, again subdivided into the same ranges of barrier parameter values.
7. Plot "KKT system condition number estimate" presents the box plot statistics of an estimate of the condition number (in logarithmic scale) for the same ranges of the barrier parameter. The estimate is obtained by a limited number of Lanczos iterations on the respective (non-)preconditioned system of the H block; a possibly remaining equality part of A is ignored in this. Computation times for the condition number are not included in the time measurement listed above.
8. Plot "preconditioning columns per KKT system" gives the box plot statistics of the number of columns $\hat{k}$ in Alg. 3 for RP and DP for the usual ranges of the barrier parameter.

In all box plots, the width of the boxes indicates the relative size of the number of instances in the group, the horizontal lines of the boxes give the values of the upper quartile, the median and the lower quartile. The upper whisker shows the largest value below upper quartile+1.5·IQR, where IQR=(upper quartile − lower quartile) is the interquartile range. The lower whisker displays the

smallest value above lower quartile−1.5·IQR. The stars show maximum and minimum value.

Computation times refer to a virtualized compute server of 40 Intel Xeon Processor (Cascadelake) cores with 600 GB RAM under Ubuntu 18.04. This virtual machine is hosted on hardware consisting of two processors Intel(R) Xeon(R) Gold 6240R CPU with 2.40GHz with 24 cores and 768 GB RAM. The code, however, is purely sequential and does not exploit any parallel computation possibilities.

4.1 Max-Cut (Instance MC, Figure 1)

The graph was randomly generated ([28], call `rudy -rnd_graph 20000 1 1` for 20000 nodes, edge density one percent, seed value 1). The semidefinite relaxation gives rise to an unconstrained problem with 20000 variables. Each variable influences one of the diagonal elements of the Laplace matrix of the graph with cost one and the task is to minimize the maximum eigenvalue of the Laplacian times the number of nodes, see [19] for the general problem description.

For graphs of this type but smaller size like 5000 or 10000 nodes the direct solver DS still seemed to perform better, so rather large sizes are needed to see some advantage of iterative methods. Other than that the relative behavior of the solvers was similar also for the smaller sizes. The jaggies within subproblem time in the second plot are due to the reduction of the model to its active part after each descent step while the model typically increases in size during null steps. During the very first iterations the bundle is tiny and DS is the best choice. Once the bundle size increases sufficiently, the iterative methods dominate. Over time, as precision requirements get higher and the choice of the bundle subspace converges, the advantage of iterative methods decreases. In the final phase of high precision the direct solver may well be more attractive again.

The plots also show that for this instance (and presumably for most instances of this random type) the performance of IT (MINRES without preconditioning) is almost as good as DP (deterministic preconditioning) while RP (randomized preconditioning) is not competitive. Note that the condition number does not grow excessively for IT in this instance. Deterministic preconditioning succeeds in keeping the condition number almost exactly at the intended value 10. For smaller values of μ, so for higher precision requirements, DP requires distinctly fewer matrix-vector multiplications, but it then also selects a large number of columns. In comparison to no preconditioning DP helps to improve stability but does not lead to significantly better computation times except maybe for the very last phase of the algorithm with high precision requirements.

4.2 Minimum Bisection (Instance BIS, Figure 2)

The semidefinite relaxation of minimum bisection is similar in nature to max-cut, but in addition to the single diagonal elements there is a variable with coefficient matrix of all ones. Furthermore, variables with sparse coefficient matrices corresponding to odd cycles in the underlying graph are added dynamically in rounds, see [14] for the general framework and also for the origin of the instance KKT_traj33 with 20006 nodes and roughly 260000 edges.

Again, after the very first iterations the iterative methods turn out to perform distinctly better in the initial phase of the algorithm. Iterative methods get less attractive as precision requirements increase. The model size is often rather small (a bit larger than the active set of about 150 columns) which is favorable for DS. Indeed, additional output information of the log file indicates that the performance of DS drops off whenever the cutting model is significantly larger than that.

While for this instance RP is better than IT, the advantage of DP over the other iterative variants is quite apparent and its superiority also increases with precision requirements and smaller μ. In fact, for DP the condition number and the number of matrix-vector multiplications decrease again for smaller μ. Possible causes might be that the active set is easier to identify correctly. Due to the reduction in matrix-vector multiplications, computation time does not increase for DP in spite of a growing number of columns in the preconditioner.

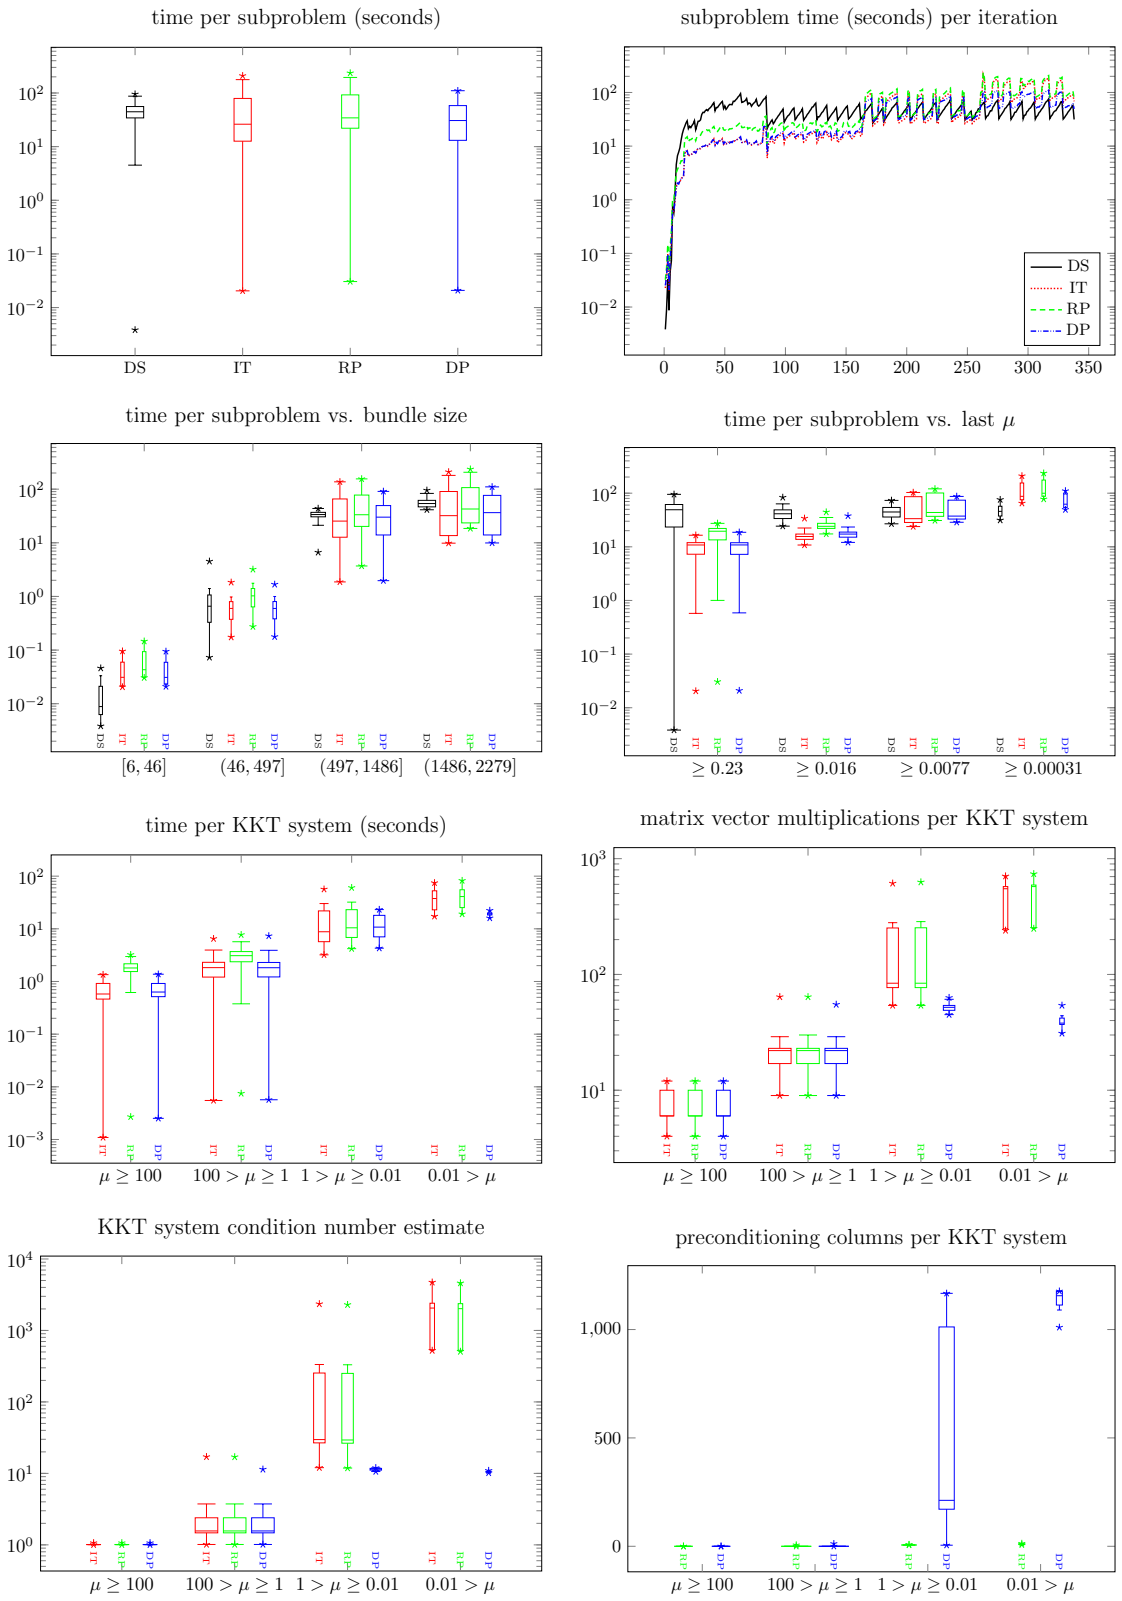

Figure 1: Instance MC, 338 subproblems, 2832 KKT systems. In order to highlight the dependence on the barrier parameter μ with corresponding precision requirements the results for KKT systems are grouped into μ-value ranges $(\infty, 100]$, $(100, 1]$, $(1, 0.01]$, $(0.01, 0)$.

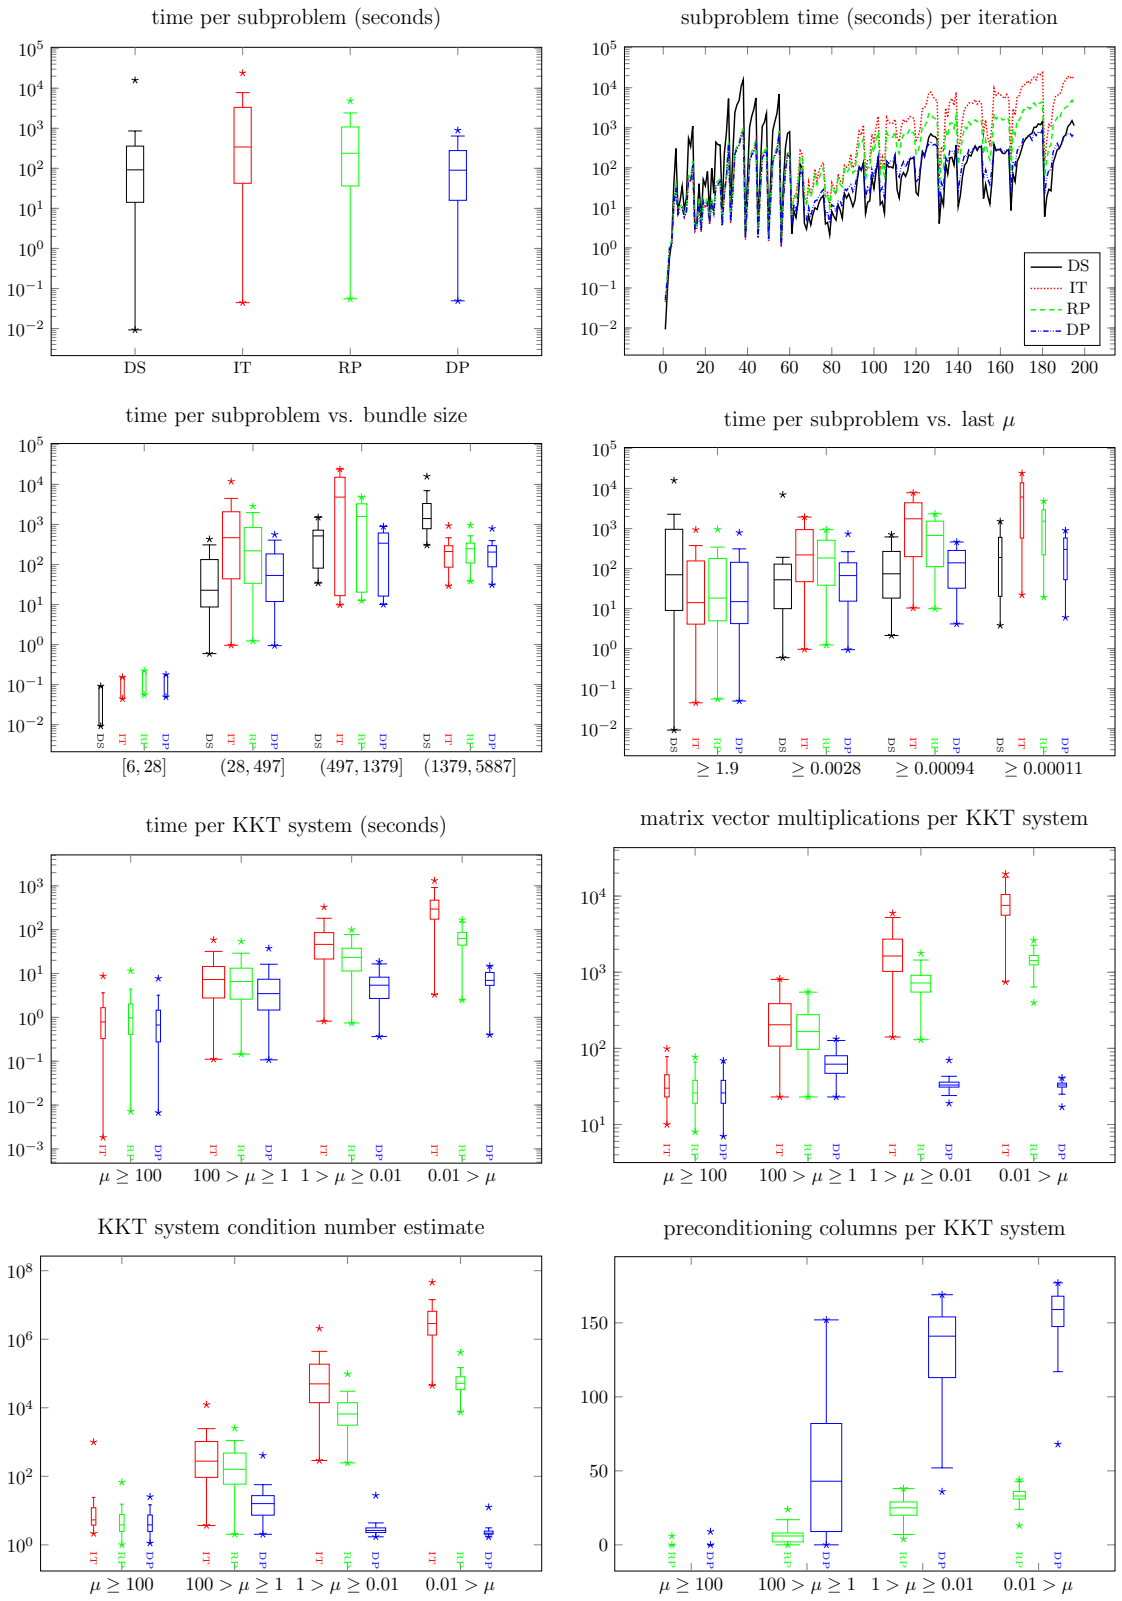

Figure 2: Instance BIS, 180 subproblems, 5380 KKT systems. In order to highlight the dependence on the barrier parameter μ with corresponding precision requirements the results for KKT systems are grouped into μ-value ranges $(\infty, 100]$, $(100, 1]$, $(1, 0.01]$, $(0.01, 0)$.

4.3 A Min-Max-Bisection Problem (Instance MMBIS, Figure 3)

This problem arose in the context of an unpublished attempt[1] to optimize vaccination rates for five population groups $N_1 \dot\cup \ldots \dot\cup N_5 = N$ in a virtual town of $n = |N|$ inhabitants. Briefly, within the town k anonymous people are assumed to be infectious. There is vaccine for at most $\underline{n}$ people. The aim is to reduce the spreading rate of the disease by vaccinating each person with the respective group's probability. The task of determining these vaccination rates motivated the following ad hoc model which would be hard to justify rigorously. In a graph $G = (N, E)$ each edge $ij = \{i, j\} \in E$ with $i \in N_{\hat\imath}$, $j \in N_{\hat\jmath}$ has a weight $\hat w_{\hat\imath\hat\jmath}$ representing the infectiousness of the typical contact for these two persons of the respective groups. It will be convenient to define the weighted Laplacians $L_{\hat\imath\hat\jmath} = \hat w_{\hat\imath\hat\jmath} \sum_{ij\in E, i\in N_{\hat\imath}, j\in N_{\hat\jmath}} (e_i - e_j)(e_i - e_j)^\top$. In this simplified approach, vaccination rates $v_{\hat\imath}, v_{\hat\jmath}$ of the node groups reduce a nominal infectiousness $\hat w_{\hat\imath\hat\jmath}$ between these groups by the factor $y_{\hat\imath\hat\jmath} \geq \max\{0, 1 - v_{\hat\imath} - v_{\hat\jmath}\}$. The spreading probability to be minimized is considered proportional to the restricted max-cut value

$$\begin{aligned}\max_{\mathcal{I}\subset N, |\mathcal{I}|=k} \sum y_{\hat\imath\hat\jmath} \hat w_{\hat\imath\hat\jmath} \cdot |\{ij \in E \colon i \in \mathcal{I} \cap N_{\hat\imath}, j \in N_{\hat\jmath} \setminus \mathcal{I}\}| = \\ = \max_{\substack{x\in\{-1,1\}^n \\ (1^\top x)^2 = (n-2k)^2}} \frac{1}{4} x^\top \Big(\sum_{\hat\imath\leq\hat\jmath} y_{\hat\imath\hat\jmath} L_{\hat\imath\hat\jmath} \Big) x,\end{aligned}$$

For determining the vaccination rates the combinatorial problem is replaced by the usual (dual) semidefinite relaxation

$$\begin{array}{rll} \text{minimize} & \frac{n}{4}\lambda_{\max}(\sum_{\hat\imath\leq\hat\jmath} y_{\hat\imath\hat\jmath} L_{\hat\imath\hat\jmath} - \operatorname{Diag}(d) - u\mathbb{1}\mathbb{1}^\top) + \mathbb{1}^\top d + (n-2k)^2 u & \\ \text{subject to} & y_{\hat\imath\hat\jmath} \geq 1 - v_{\hat\imath} - v_{\hat\jmath}, & \hat\imath \leq \hat\jmath, \\ & \sum_{\hat\imath} |N_{\hat\imath}| v_{\hat\imath} = \underline{n}, & \\ & d \in \mathbb{R}^n, u \in \mathbb{R}, y \geq 0, v \geq 0. & \end{array}$$

In this case the resulting KKT system also has an equality and several inequality constraints in the block A. Preconditioning results are presented for the KKT systems of an instance with $n = 12600$ inhabitants splitting into groups of sizes $5770, 6000, 600, 30, 200$, with $k = 126$ infectious persons and $\underline{n} = 1260$ available vaccinations.

In the actual computations the bundle size grows surprisingly fast. This not only entails enormous memory requirements but also excessive computation times for DS; indeed, computations of DS may exceed those of DP by a factor of 70. In consequence comparative results can only be reported for a very limited number of subproblem evaluations. In particular, the precision requirements remain rather moderate throughout these iterations. Still, the same initial behavior can be observed as for the previous two instances. For very small bundle sizes DS is best. Once the bundle size grows, the iterative methods take over. Among the iterative solvers RP is better than IT, but DP is the method of choice. It succeeds in tightly controlling the condition number by selecting rather few columns. With this DP requires the fewest matrix vector multiplications which seems to pay of quickly on this instance.

5 Conclusions

In search for efficient low rank preconditioning techniques for the iterative solution of the internal KKT system of the quadratic bundle subproblem two subspace selection heuristics — a randomized and a deterministic variant — were proposed. For the randomized approach the results are ambivalent in theory and in practice; obtaining a good subspace this way seems to be difficult and the cost of exploratory matrix-vector multiplications quickly dominates. In contrast, the deterministic subspace selection approach allows to control the condition number (and with it the number of matrix vector multiplications) at a desired level without the need to tune any parameters in

[1] together with B. Filipecki (TU Chemnitz), S. Heyder (TU Ilmenau), Th. Hotz (TU Ilmenau) within BMBF-project grant 05M18OCA.

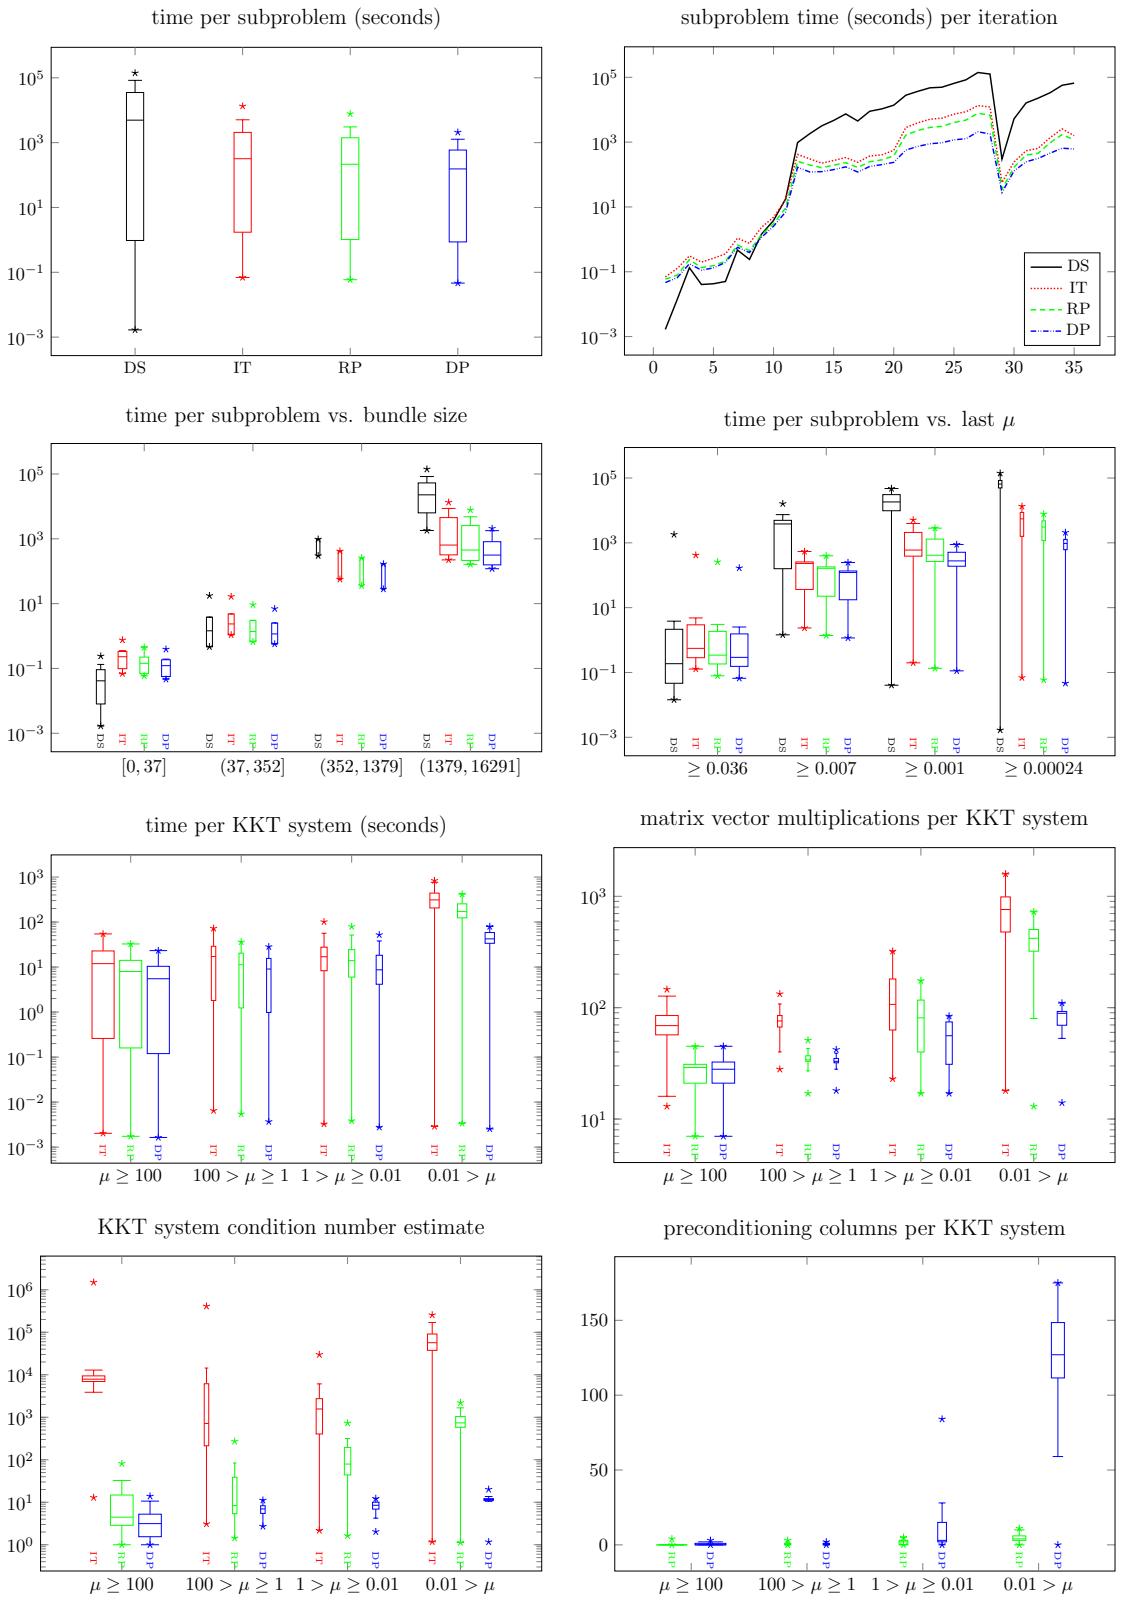

Figure 3: Instance MMBIS, 28 subproblems, 642 KKT systems. In order to highlight the dependence on the barrier parameter μ with corresponding precision requirements the results for KKT systems are grouped into μ-value ranges $(\infty, 100]$, $(100, 1]$, $(1, 0.01]$, $(0.01, 0)$.

theory as well as on the test instances. On these instances, for low precision requirements (large barrier parameter) the selected subspace is neglectably small. For high precision requirements (small barrier parameter) the subspace grows to the active model subspace. If the bundle size is close to this active dimension, the work in forming the preconditioner may be comparable to forming the Schur complement for the direct solver. Still, for large scale instances the deterministically preconditioned iterative approach seems to be preferable.

Conceivably it is possible to profit in ConicBundle from the advantages of the deterministic iterative and the direct solver by switching dynamically between both. The current experiments relied on a predictor-corrector approach that was tuned for the direct solver. In view of the properties of the iterative approach it may well be worth to devise a different path following strategy for the iterative approach, in particular for the initial phase of the interior point method when the barrier parameter is still comparatively large and the work invested in forming the preconditioner is still neglectable. Similar ideas should be applicable to interior point solvers for solving convex quadratic problems with low rank structure.

Acknowledgments I have profited a lot from discussions with many colleagues, in part years back. In particular I have to thank K.-C. Toh as well as my colleagues O. Ernst, R. Herzog, A. Pichler and M. Stoll in Chemnitz. Much of the preparatory restructuring of ConicBundle was done during my sabbatical at the University of Klagenfurt, thank you to F. Rendl and A. Wiegele for making this possible. The support of research grants 05M18OCA of the German Federal Ministry of Education and Research and the DFG CRC 1410 is gratefully acknowledged.

References

[1] D. Achlioptas. Database friendly random projections. In *Proc 20th ACM Symp Principles of Database Systems*, Santa Barabara, CA, 2001. 274–281.

[2] F. Alizadeh, J.-P. A. Haeberly, and M. L. Overton. Primal-dual interior-point methods for semidefinite programming: Convergence rates, stability and numerical results. *SIAM J. Optim.*, 8(3):746–768, Aug. 1998.

[3] M. F. Anjos and J. B. Lasserre, editors. *Handbook of Semidefinite, Conic and Polynomial Optimization*, volume 166 of *International Series in Operations Research & Management Science*. Springer, 2012.

[4] S. Benson, Y. Ye, and X. Zhang. Solving large-scale sparse semidefinite programs for combinatorial optimization. *SIAM J. Optim.*, 10(2):443–461, 2000.

[5] J. F. Bonnans, J. C. Gilbert, C. Lemaréchal, and C. A. Sagastizábal. *Numerical Optimization*. Springer, 2nd edition, 2006.

[6] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004. Reprinted 2007 with corrections.

[7] S. Burer and R. D. Monteiro. A nonlinear programming algorithm for solving semidefinite programs via low-rank factorization. *Math. Programming*, 94:329–357, 2003.

[8] S. Dasgupta and A. Gupta. An elementary proof of a theorem of Johnson and Lindenstrauss. *Random Structures and Algorithms*, 22(1):60–65, 2002. https://doi.org/10.1002/rsa.10073.

[9] L. Ding and B. Grimmer. Revisit of spectral bundle methods: Primal-dual (sub)linear convergence rates, 2020. arXiv:2008.07067.

[10] H. C. Elman, D. J. Silvester, and A. J. Wathen. *Finite Elements and Fast Iterative Solvers: with Applications in Incompressible Fluid Dynamics*. Oxford University Press, 2005. Reprinted 2006.

[11] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. ACM*, 42:1115–1145, 1995.

[12] S. Habibi, A. Kavand, M. Kocvara, and M. Stingl. Barrier and penalty methods for low-rank semidefinite programming with application to truss topology design, 2021. arXiv:2105.08529.

[13] N. Halko, P. G. Martinsson, and J. A. Tropp. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM Review*, 53(2):217–288, 2011.

[14] C. Helmberg. A cutting plane algorithm for large scale semidefinite relaxations. In M. Grötschel, editor, *The Sharpest Cut*, MPS-SIAM Series on Optimization, pages 233–256. SIAM/MPS, 2004.

[15] C. Helmberg. *ConicBundle v1.a.2*. Fakultät für Mathematik, Technische Universität Chemnitz, 2021. `http://www.tu-chemnitz.de/~helmberg/ConicBundle`.

[16] C. Helmberg and K. C. Kiwiel. A spectral bundle method with bounds. *Math. Programming*, 93(2):173–194, 2002.

[17] C. Helmberg, M. L. Overton, and F. Rendl. The spectral bundle method with second-order information. *Optimization Methods and Software*, 29(4):855–876, July 2014.

[18] C. Helmberg and A. Pichler. Dynamic scaling and submodel selection in bundle methods for convex optimization. Preprint 2017, Fakultät für Mathematik, Technische Universität Chemnitz, D-09107 Chemnitz, Germany, Aug. 2017.

[19] C. Helmberg and F. Rendl. A spectral bundle method for semidefinite programming. *SIAM J. Optim.*, 10(3):673–696, 2000.

[20] D. Henrion, M. Korda, and H. B. Lasserre. *The Moment-SOS Hierarchy*, volume 4 of *Series on Optimization and Its Applications*. World Scientific, 2020. ISBN: 978-1-78634-853-1.

[21] N. J. Higham and T. Mary. A new preconditioner that exploits low-rank approximations to factorization error. *SIAM J. Sci. Comput.*, 41(1):A59–A82, 2019.

[22] J.-B. Hiriart-Urruty and C. Lemaréchal. *Convex Analysis and Minimization Algorithms I*, volume 305 of *Grundlehren der mathematischen Wissenschaften*. Springer, Berlin, Heidelberg, 1993.

[23] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 1985.

[24] S. Kim, M. Kojima, M. Mevissen, and M. Yamashita. Exploiting sparsity in linear and nonlinear matrix inequalities via positive semidefinite matrix completion. *Math. Programming*, 129:33–68, 2011.

[25] M. Kocvara and M. Stingl. On the solution of large-scale sdp problems by the modified barrier method using iterative solvers. *Math. Programming*, 109(2-3):413–444, 2007.

[26] Y. Nesterov and M. J. Todd. Primal-dual interior-point methods for self-scaled cones. *SIAM J. Optim.*, 8:324–364, 1998.

[27] C. H. Papadimitriou, P. Raghavan, H. Tamaki, and S. Vempala. Latent semantic indexing: A probabilistic analysis. *Journal of Computer and System Sciences*, 61:217–235, 2000. doi:10.1006/jcss.2000.1711.

[28] G. Rinaldi. rudy, 1993. A rudimental graph generator by JRT, available at `https://www-user.tu-chemnitz.de/~helmberg/rudy.tar.gz`.

[29] R. T. Rockafellar. *Convex Analysis*. Princeton University Press, Princeton, New Jersey, 1970.

[30] M. J. Todd, K. C. Toh, and R. H. Tütüncü. On the Nesterov-Todd direction in semidefinite programming. *SIAM J. Optim.*, 8(3):769–796, Aug. 1998.

[31] K.-C. Toh. An inexact primal-dual path following algorithm for convex quadratic SDP. *Math. Programming*, 112(1):221–254, 2008.

[32] H. Wolkowicz, R. Saigal, and L. Vandenberghe, editors. *Handbook of Semidefinite Programming*, volume 27 of *International Series in Operations Research and Management Science*. Kluwer Academic Publishers, Boston/Dordrecht/London, 2000.

[33] R. Y. Zhang and J. Lavaei. Modified interior-point method for large-and-sparse low-rank semidefinite programs. In *2017 IEE 56th Conference on Decision and Control (CDC)*, pages 5640–5647, Melbourne, Australia, December 12–15 2017.

A Tables

For each box plot of figures 1–3 for the three instances MC, BIS and MMBIS the following tables list the number of instances and the values of the parameters minimum, lower quartile (Q_1), median, upper quartile (Q_3), maximum. For each of the three instances an additional table gives the statistics on the Euclidean norm of the resulting residual of (6) achieved by the respective solver for the KKT-systems grouped by the usual value ranges of the barrier parameter.

A.1 Max-Cut (Instance MC, Figure 1)

Time per subproblem in seconds (338 instances):

solver	min	Q_1	median	Q_3	max
DS	0.003832	34.3212	44.6572	55.9219	95.1265
IT	0.020491	12.6342	26.1729	79.1747	209.58
RP	0.030548	22.0347	34.3771	92.207	235.653
DP	0.020898	13.0997	30.6566	58.1994	110.33

Time per subproblem in seconds vs. ranges of bundle sizes:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[6, 46]$	5	DS	0.003832	0.006275	0.0088455	0.021121	0.045992
		IT	0.020491	0.021558	0.031059	0.059577	0.095407
		RP	0.030548	0.0314005	0.043048	0.093556	0.146144
		DP	0.020898	0.0230755	0.0309575	0.0592245	0.094453
$[56, 497]$	5	DS	0.073058	0.329229	0.657365	1.0668	4.51608
		IT	0.173868	0.374044	0.597913	0.800166	1.83209
		RP	0.273307	0.637968	1.02498	1.4059	3.19528
		DP	0.177225	0.382114	0.59731	0.80289	1.6773
$[596, 1486]$	133	DS	6.62451	30.4929	33.6252	36.9311	43.6407
		IT	1.86562	12.7032	25.3882	65.6882	136.935
		RP	3.69818	20.2136	33.3548	77.6126	155.415
		DP	1.96853	13.9917	30.0676	49.426	90.8324
$[1541, 2279]$	195	DS	41.2984	46.7501	53.9483	61.8286	95.1265
		IT	9.84382	13.6284	32.0618	90.4042	209.58
		RP	18.53	23.4963	42.6607	106.814	235.653
		DP	9.9354	13.9977	36.494	76.5581	110.33

Time per subproblem in seconds vs. last barrier parameter μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[36, 0.23]$	85	DS	0.003832	23.4027	49.1143	61.5018	95.1265
		IT	0.020491	7.2764	10.8231	11.9651	16.4467
		RP	0.030548	13.4739	19.5347	22.0031	27.3839
		DP	0.020898	7.2377	10.8421	11.9242	18.6599
$[0.23, 0.016]$	85	DS	24.2787	33.6785	41.3086	48.8582	83.315
		IT	10.8004	13.7717	15.5724	17.2952	33.9445
		RP	17.3893	22.0349	24.1102	27.3977	44.4086
		DP	12.0901	15.1287	17.2843	18.8158	37.7594
$[0.016, 0.0077]$	85	DS	26.8515	35.9349	44.6111	54.1812	73.3521
		IT	23.8603	28.455	33.4883	86.0521	103.067
		RP	30.9349	37.0011	43.6655	100.814	120.387
		DP	28.6909	32.9195	37.3972	74.0808	87.0083
$[0.0077, 0.00031]$	83	DS	31.4011	37.4751	45.8135	57.0963	75.7966
		IT	64.8852	75.3411	87.3138	154.968	209.58
		RP	77.5913	87.3499	99.7462	175.749	235.653
		DP	49.3655	53.8412	62.3306	96.8023	110.33

Time per KKT system in seconds grouped by value ranges of the barrier parameter μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[4.5e+02, 9.9e+02]$	676	IT	0.001092	0.463337	0.578703	0.917693	1.34934
		RP	0.002696	1.54367	1.80541	2.1684	3.22664
		DP	0.00252	0.515255	0.630482	0.916957	1.37408
$[1.4, 51]$	1490	IT	0.005506	1.2121	1.83286	2.31178	6.49419
		RP	0.007486	2.36829	3.09237	3.69645	7.72418
		DP	0.005681	1.22031	1.82556	2.29723	7.32683
$[0.016, 0.31]$	510	IT	3.20468	5.69856	8.79662	21.8884	56.4563
		RP	4.19829	6.85746	10.4541	23.0927	60.4059
		DP	4.28737	7.03695	10.7758	18.0144	23.1913
$[0.00031, 0.0093]$	156	IT	17.2603	22.94	37.6166	52.734	74.6325
		RP	19.0854	25.2808	40.9692	55.18	82.3144
		DP	16.0158	18.3264	19.1491	19.8728	22.2577

Number of matrix vector multiplications per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[4.5e+02, 9.9e+02]$	676	IT	4	6	6	10	12
		RP	4	6	6	10	12
		DP	4	6	6	10	12
$[1.4, 51]$	1490	IT	9	17	22	23	64
		RP	9	17	22	23	64
		DP	9	17	22	23	55
$[0.016, 0.31]$	510	IT	54	77	84	252	612
		RP	54	77	84	253	627
		DP	45	49	52	54	63
$[0.00031, 0.0093]$	156	IT	240	248	550	574	706
		RP	249	253	574	592	739
		DP	31	37	38	42	54

Condition number estimate of the KKT systems grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[4.5e+02, 9.9e+02]$	676	IT	1.003	1.009	1.009	1.01	1.065
		RP	1.003	1.009	1.009	1.01	1.065
		DP	1.003	1.009	1.009	1.01	1.065
$[1.4, 51]$	1490	IT	1.013	1.473	1.566	2.393	17.06
		RP	1.013	1.473	1.566	2.393	16.98
		DP	1.013	1.473	1.566	2.393	11.38
$[0.016, 0.31]$	510	IT	12.03	26.72	29.72	253.7	2349
		RP	11.87	26.44	29.33	250.7	2289
		DP	10.62	11.11	11.46	11.75	11.96
$[0.00031, 0.0093]$	156	IT	523.5	536.4	2057	2404	4711
		RP	507.1	527.3	2017	2371	4599
		DP	10.17	10.43	10.56	10.66	10.93

Number of preconditioning columns per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[4.5e+02, 9.9e+02]$	676	RP	0	0	0	0	0
		DP	0	0	0	0	0
$[1.4, 51]$	1490	RP	0	0	0	0	6
		DP	0	0	0	0	12
$[0.016, 0.31]$	510	RP	6	6	6	6	8
		DP	5	171	212	1012	1167
$[0.00031, 0.0093]$	156	RP	7	8	10	10	13
		DP	1010	1113	1156	1175	1177

Euclidean norm of the residual of (6) per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[4.5e+02, 9.9e+02]$	676	IT	6.4e-10	1.4e-09	2.6e-09	5.2e-09	1e-06
		RP	6.4e-10	1.4e-09	2.6e-09	5.2e-09	1e-06
		DP	6.4e-10	1.4e-09	2.6e-09	5.2e-09	1e-06
$[1.4, 51]$	1490	IT	3.6e-10	6.7e-09	5.2e-08	3.9e-07	1e-06
		RP	3.6e-10	6.4e-09	5.2e-08	3.9e-07	1.1e-06
		DP	3.6e-10	6.7e-09	5.2e-08	3.9e-07	1e-06
$[0.016, 0.31]$	510	IT	8.6e-09	7e-07	8.2e-07	9.2e-07	1e-06
		RP	9.7e-09	7.3e-07	9.2e-07	1.1e-06	2.4e-06
		DP	5.2e-09	5.4e-07	6.8e-07	7.9e-07	1.2e-06
$[0.00031, 0.0093]$	156	IT	9.7e-09	9.6e-07	9.8e-07	9.9e-07	1e-06
		RP	1.9e-08	1.4e-06	2.2e-06	4.2e-06	6.2e-06
		DP	8.6e-09	5.8e-07	7.1e-07	8.1e-07	1e-06

A.2 Minimum Bisection (Instance BIS, Figure 2)

Time per subproblem in seconds (195 instances):

solver	min	Q_1	median	Q_3	max
DS	0.00923	14.1723	91.7467	359.175	15841.2
IT	0.044218	42.3879	340.967	3329.25	24016.1
RP	0.055209	36.3358	238.703	1079.8	4866.25
DP	0.049162	15.9154	89.9136	277.508	890.086

Time per subproblem in seconds vs. ranges of bundle sizes:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[6, 28]$	2	DS	0.00923	–	0.00923	–	0.091545
		IT	0.044218	–	0.044218	–	0.153935
		RP	0.055209	–	0.055209	–	0.222727
		DP	0.049162	–	0.049162	–	0.177788
$[67, 497]$	128	DS	0.594552	8.69331	22.7682	132.809	428.513
		IT	0.960953	43.8529	466.297	2076.04	11867
		RP	1.23022	33.9205	219.14	835.573	2856.49
		DP	0.941817	11.8932	53.3373	183.946	560.093
$[529, 1379]$	40	DS	34.3749	81.2278	517.036	720.019	1505.5
		IT	9.82356	16.6676	4819.87	15092.7	24016.1
		RP	12.5897	20.357	1586.65	3277.46	4866.25
		DP	10.0625	16.2665	340.397	612.94	890.086
$[1597, 5887]$	25	DS	302.944	783.079	1403.96	3335.16	15841.2
		IT	29.3698	85.8307	211.467	294.115	932.896
		RP	38.2829	108.971	248.069	341.147	955.341
		DP	31.1555	87.6845	204.633	295.527	789.044

Time per subproblem in seconds vs. last barrier parameter μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1e+03, 1.9]$	49	DS	0.00923	8.97908	69.8681	955.047	15841.2
		IT	0.044218	4.08472	14.083	154.701	932.896
		RP	0.055209	4.93747	18.2564	178.33	955.341
		DP	0.049162	4.21385	14.8836	143.015	789.044
$[1.4, 0.0028]$	49	DS	0.594552	9.95126	52.3298	128.666	6981.47
		IT	0.960953	46.9621	218.946	937.902	1927.16
		RP	1.23022	38.2691	183.586	508.265	933.538
		DP	0.941817	15.3137	66.5916	139.005	734.869
$[0.0028, 0.00094]$	49	DS	2.12836	18.2489	73.8103	267.252	705.689
		IT	10.3255	198.11	1744.97	4383.54	7789.69
		RP	9.92439	111.622	670.068	1522.76	2315.44
		DP	4.13901	32.1878	138.761	282.768	459.79
$[0.00094, 0.00011]$	48	DS	3.77535	20.2483	189.134	599.782	1505.5
		IT	21.8778	575.065	6101.11	13836.1	24016.1
		RP	19.1305	219.14	1524.83	2942.66	4866.25
		DP	5.98724	52.4419	299.974	580.588	890.086

Time per KKT system in seconds grouped by value ranges of the barrier parameter μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1.1e+02, 2.6e+04]$	824	IT	0.001823	0.328123	0.786603	1.66599	8.7638
		RP	0.007159	0.410637	0.980222	2.023	11.5388
		DP	0.006695	0.275799	0.669313	1.45762	7.73375
$[1, 99]$	2262	IT	0.110818	2.7715	7.29372	14.4584	58.5776
		RP	0.145002	2.60334	6.57376	13.2324	54.0312
		DP	0.10742	1.46962	3.4827	7.40111	37.485
$[0.01, 1]$	1991	IT	0.817783	21.4256	46.2355	85.6325	327.564
		RP	0.744242	11.4276	23.3757	37.8191	100.053
		DP	0.365208	2.70049	5.41858	8.25427	18.679
$[0.00011, 0.01]$	1103	IT	3.30796	173.717	295.614	470.923	1316.81
		RP	2.48871	44.4168	62.784	86.7278	168.322
		DP	0.40273	5.37352	6.99157	10.5791	14.8519

Number of matrix vector multiplications per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1.1e+02, 2.6e+04]$	824	IT	10	23	30	45	99
		RP	8	19	26	38	77
		DP	7	19	26	38	69
$[1, 99]$	2262	IT	23	107	204	386.5	820
		RP	23	97	166	277	555
		DP	23	47	62	80	133
$[0.01, 1]$	1991	IT	141	1024.5	1632	2719.5	5960
		RP	130	548.5	722.5	910.5	1774
		DP	19	31	33	36	70
$[0.00011, 0.01]$	1103	IT	750	5607.5	7552.5	10475	19571
		RP	399	1249	1419.5	1657	2620
		DP	17	31	33	35	41

Condition number estimate of the KKT systems grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1.1e+02, 2.6e+04]$	824	IT	2.15	3.773	5.36	12.1	997.8
		RP	1.008	2.465	3.84	7.696	67.38
		DP	1.122	2.469	3.834	7.442	25.34
$[1, 99]$	2262	IT	3.681	93.22	277.5	1041	1.234e+04
		RP	2.035	59.19	160	477	2544
		DP	2.035	7.361	15.99	27.31	407.7
$[0.01, 1]$	1991	IT	289.5	1.409e+04	4.986e+04	1.869e+05	2.075e+06
		RP	246.8	3107	6609	1.405e+04	9.741e+04
		DP	1.717	2.29	2.632	3.128	27.66
$[0.00011, 0.01]$	1103	IT	4.501e+04	1.324e+06	2.869e+06	6.53e+06	4.576e+07
		RP	7554	3.375e+04	5.161e+04	8.013e+04	4.099e+05
		DP	1.703	2.048	2.185	2.522	12.55

Number of preconditioning columns per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1.1e+02, 2.6e+04]$	824	RP	0	0	0	0	6
		DP	0	0	0	0	9
$[1, 99]$	2262	RP	0	2	6	8	24
		DP	0	9	43	82	152
$[0.01, 1]$	1991	RP	4	20	25	29	38
		DP	36	113	141	154	169
$[0.00011, 0.01]$	1103	RP	13	31	33	36	44
		DP	68	147.5	159	168	177

Euclidean norm of the residual of (6) per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1.1e+02, 2.6e+04]$	824	IT	9.8e-13	9.1e-09	3.2e-07	5.6e-07	1e-06
		RP	1.4e-12	6.9e-09	1.9e-07	3.5e-07	1.7e-06
		DP	1.3e-12	7e-09	1.9e-07	3.5e-07	1.1e-06
$[1, 99]$	2262	IT	2.9e-10	4.7e-09	6.7e-09	5e-07	1e-06
		RP	3.1e-10	4e-09	5.7e-09	3.3e-07	3.2e-06
		DP	3.8e-10	5e-09	7.4e-09	3.5e-07	2.8e-06
$[0.01, 1]$	1991	IT	1.2e-09	5e-09	9e-09	1.9e-08	1e-06
		RP	1.6e-09	6.4e-09	1.2e-08	2.4e-08	2.2e-05
		DP	6.6e-10	5.2e-09	9.4e-09	2e-08	7.2e-06
$[0.00011, 0.01]$	1103	IT	4.3e-09	4.9e-08	1.1e-07	3.3e-07	2.9e-06
		RP	3.5e-09	3.3e-08	5.7e-08	9.5e-08	7.3e-05
		DP	2.2e-09	3.4e-08	6e-08	1e-07	5.3e-06

A.3 Min-Max Bisection (Instance MMBIS, Figure 3

Time per subproblem in seconds (35 instances):

solver	min	Q_1	median	Q_3	max
DS	0.001669	0.958311	4942	35185.7	140651
IT	0.06925	1.7246	319.124	2069.75	13368.7
RP	0.059093	1.02453	215.038	1412.83	7763.78
DP	0.046371	0.864647	154.628	590.86	2083.68

Time per subproblem in seconds vs. ranges of bundle sizes:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[0, 37]$	7	DS	0.001669	0.008018	0.0415445	0.091635	0.240095
		IT	0.06925	0.098375	0.229849	0.332737	0.752143
		RP	0.059093	0.069004	0.144198	0.224022	0.449136
		DP	0.046371	0.056267	0.121948	0.185404	0.391171
$[56, 352]$	4	DS	0.46651	0.46651	1.45011	3.83018	17.636
		IT	1.08647	1.08647	2.36273	4.81555	16.5215
		RP	0.668257	0.668257	1.3808	3.04164	9.26137
		DP	0.567064	0.567064	1.16223	2.53117	6.96995
$[1327, 1379]$	2	DS	300.616	–	300.616	–	969.539
		IT	56.453	–	56.453	–	416.7
		RP	35.8027	–	35.8027	–	257.055
		DP	28.0559	–	28.0559	–	167.41
$[2702, 16291]$	22	DS	1807.41	6318.51	22730.5	53047.4	140651
		IT	223.837	319.124	641.539	4509.64	13368.7
		RP	163.342	215.038	450.345	2558.71	7763.78
		DP	119.399	157.057	315.234	814.558	2083.68

Time per subproblem in seconds vs. last barrier parameter μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1.4, 0.036]$	9	DS	0.014367	0.0464135	0.186595	2.14835	1807.41
		IT	0.1275	0.285552	0.553999	2.95101	416.7
		RP	0.078915	0.183008	0.342684	1.85495	257.055
		DP	0.066163	0.154237	0.292596	1.54912	167.41
$[0.035, 0.007]$	9	DS	1.45011	159.126	3825.74	4942	16156.8
		IT	2.36273	36.4873	230.856	257.261	539.789
		RP	1.3808	22.5321	163.699	181.259	397.781
		DP	1.16223	17.5129	121.09	135.255	246.92
$[0.0059, 0.001]$	9	DS	0.040436	9770.54	18267.6	30691.4	47352.7
		IT	0.198213	388.296	598.446	2093.91	5072.23
		RP	0.134192	265.913	414.119	1302.96	2838.2
		DP	0.11221	189.641	276.743	511.322	890.336
$[0.00096, 0.00024]$	8	DS	0.001669	49409.5	64909.4	83523.5	140651
		IT	0.06925	1582.35	5497.04	8664.78	13368.7
		RP	0.059093	1175.55	3060.07	4787.77	7763.78
		DP	0.046371	614.286	954.903	1273.72	2083.68

Time per KKT system in seconds grouped by value ranges of the barrier parameter μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1e+02, 1e+03]$	295	IT	0.002025	0.259585	11.9083	22.8061	54.4588
		RP	0.001737	0.159694	7.95722	14.0881	32.6083
		DP	0.001638	0.119556	5.4995	10.3867	23.2037
$[1, 99]$	146	IT	0.00647	1.8041	17.2027	28.9793	72.3436
		RP	0.005452	1.23622	11.287	20.3489	36.1462
		DP	0.003676	0.98141	9.03463	15.5483	28.2709
$[0.012, 1]$	159	IT	0.003261	8.2878	16.9553	27.6698	101.194
		RP	0.003848	5.98822	13.8806	24.4045	79.1858
		DP	0.002763	4.14727	8.67571	18.2895	52.1301
$[0.00024, 0.0086]$	182	IT	0.002879	205.544	310.59	439.276	833.569
		RP	0.003367	123.886	170.845	253.51	417.102
		DP	0.00254	33.7196	42.3111	58.5254	79.5814

Number of matrix vector multiplications per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1e+02, 1e+03]$	295	IT	13	57	69	85	146
		RP	7	21	29	31	45
		DP	7	21	28	32.5	45
$[1, 99]$	146	IT	28	67	76	85	133
		RP	17	33	34	37	51
		DP	18	32	33	35	42
$[0.012, 1]$	159	IT	23	63	107	181	321
		RP	17	40	81	117	175
		DP	17	31	56	74.5	84
$[0.00024, 0.0086]$	182	IT	18	478	761	987.5	1590
		RP	13	321.5	420	503	726
		DP	14	69.5	89	92.5	110

Condition number estimate of the KKT systems grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1e+02, 1e+03]$	295	IT	12.93	6950	7879	9417	1.496e+06
		RP	1.001	2.869	4.467	14.82	81.57
		DP	1.001	1.549	3.146	5.245	13.87
$[1, 99]$	146	IT	3.102	213.5	714.8	6152	4.1e+05
		RP	1.434	5.384	8.388	38.71	271.3
		DP	2.706	5.471	6.97	8.305	11.16
$[0.012, 1]$	159	IT	2.177	404.9	1566	2739	2.984e+04
		RP	1.639	43.97	79.24	195.1	730.4
		DP	2.031	6.9	8.502	10.03	12.24
$[0.00024, 0.0086]$	182	IT	1.179	3.755e+04	5.674e+04	9.099e+04	2.553e+05
		RP	1.131	580.6	737.5	1039	2200
		DP	1.169	11.08	11.61	12.17	20.08

Number of preconditioning columns per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1e+02, 1e+03]$	295	RP	0	0	0	0	4
		DP	0	0	0	1	3
$[1, 99]$	146	RP	0	0	0	1	3
		DP	0	0	0	1	2
$[0.012, 1]$	159	RP	0	0	2	3	5
		DP	0	2	3	15	84
$[0.00024, 0.0086]$	182	RP	0	3	4	6	11
		DP	0	111.5	127	148.5	175

Euclidean norm of the residual of (6) per KKT system grouped by value ranges of μ:

μ-range	#	solver	min	Q_1	median	Q_3	max
$[1e+02, 1e+03]$	295	IT	1e-09	7.8e-08	2.5e-07	5e-07	9.8e-07
		RP	2.7e-11	3.6e-09	1e-08	5e-08	5.3e-07
		DP	3.4e-11	3.8e-09	1.2e-08	3.8e-08	3.9e-07
$[1, 99]$	146	IT	3.3e-09	3.4e-08	3.3e-07	6.5e-07	1e-06
		RP	1.2e-12	2.3e-08	9.1e-08	1.8e-07	1.1e-06
		DP	2.2e-10	1.7e-08	7.6e-08	2.5e-07	1.1e-06
$[0.012, 1]$	159	IT	6.5e-10	1.3e-08	1.4e-07	4.9e-07	9.8e-07
		RP	5.4e-10	8.5e-09	8.8e-08	1.7e-07	7.7e-07
		DP	2e-10	1.1e-08	1.2e-07	2.9e-07	8.9e-07
$[0.00024, 0.0086]$	182	IT	3.5e-09	3.4e-08	7.8e-08	6.5e-07	1.3e-06
		RP	6.5e-10	1.2e-08	2.3e-08	8.5e-08	7.7e-07
		DP	7.8e-10	1.5e-08	3.1e-08	1e-07	9e-07